DOCUMENT RESUME

ED 151 704 CG 012 321

AUTHOR	Underwood, Benton J.; Malmi, Robert A.
TITLE	Studies on the Acquisition of Temporal Codes for Words Within a List.
SPONS AGENCY	Office of Naval Research, Arlington, Va. Personnel and Training Research Programs Office.
BUREAU NO	RR-042-06
PUB DATE	Jul 77
CONTRACT	N00014-76-C-0270
NOTE	64p.
EDRS PRICE	MF-$0.83 HC-$3.50 Plus Postage.
DESCRIPTORS	Adults; Cues; Experimental Psychology; *Learning Processes; *Memory; *Recall (Psychological); Research Projects; *Retention; Serial Learning; *Time Factors (Learning)
IDENTIFIERS	*Recency Principle; *Temporal Codes

ABSTRACT

Several different issues in the temporal coding of words were subjected to experimental analysis. Two experiments evaluated three response measures (recency judgments, position judgments, lag judgments) used to index temporal coding. Lag judgments were found to be of little use; subjects could make valid position and recency judgments without being able to make valid lag judgments. Practicing lag judgments produced heavy positive transfer to the other two measures. Experiment 3 showed that correct recency judgments were a direct linear function of rate of presentation through three seconds. Experiment 4 demonstrated that recency learning and two-category classification learning were substantially correlated, but a direct test (Experiment 5) indicated that the two-category classification task cannot serve as a paradigm for recency learning. Experiment 5 also showed that word frequency had no influence on either recency learning or two-category classification learning. Experiment 6 suggested that a recency principle may govern knowledge of temporal order for very short intervals of time. (Author)

****
* Reproductions supplied by EDRS are the best that can be made *
* from the original document. *
****

ED151704

STUDIES ON THE ACQUISITION OF TEMPORAL CODES FOR WORDS WITHIN A LIST

U.S. DEPARTMENT OF HEALTH, EDUCATION & WELFARE
NATIONAL INSTITUTE OF EDUCATION

THIS DOCUMENT HAS BEEN REPRODUCED EXACTLY AS RECEIVED FROM THE PERSON OR ORGANIZATION ORIGINATING IT. POINTS OF VIEW OR OPINIONS STATED DO NOT NECESSARILY REPRESENT OFFICIAL NATIONAL INSTITUTE OF EDUCATION POSITION OR POLICY

Benton J. Underwood & Robert A. Malmi

Northwestern University

July, 1977

Project RR 042-06

Sponsored by
Personnel & Training Research Programs
Psychological Sciences Division
Office of Naval Research
Arlington, Virginia
Contract No. N00014-76-C-0270

Approved for public release; distribution unlimited.
Reproduction in whole or in part is permitted
for any purpose of the United States Government.

GG 012321

Studies on the Acquisition of Temporal Codes for Words Within A List

Benton J. Underwood & Robert A. Malmi

When a relatively discreet event occurs, the memory for that event may carry information which reflects the point in time at which the memory was established. Thus, it is said that memories may be temporally coded. Many of our laboratory tasks differ in the demands placed on temporal coding. Serial learning and free-recall learning differ in that in the former, recall must be ordered to correspond to presentation order, whereas this is not required in the latter type of learning. Still, free-recall learning requires a distinction between words presented for study and other words not presented, a distinction which in recognition tests is identified explicitly with the temporal terms "old" and "new." Retroactive and proactive inhibition between two lists may result in part from the lack of (or loss of) information concerning the order of the lists in learning.

Certain issues and problems associated with attempts to understand the development of temporal codes for memories were identified in a previous report (Underwood, 1977). The studies to be presented here represent efforts to clarify certain of the issues, and also to make some theoretical tests of mechanisms presumed to be involved in temporal coding. As implied above, temporal coding may be studied for successively presented lists of words (temporal differentiation among lists), or for successively presented words within a list. Our experiments were concerned exclusively with the temporal knowledge for individual words within a list. The purposes and backgrounds of the six experiments were somewhat diverse. Therefore, each will be given a separate introduction.

Experiment I

Three response measures have been used to index temporal coding for words within a list. The description of each assumes that the subject has been shown a list of words singly for study. The response measures differ because different questions are asked of the subjects. The first response measure to be described will be called position judgments. On the test, words from the list are shown the subjects and they are asked to identify the position held by each on the study trial. The exact position may be requested, or, more grossly, the subjects may be asked to identify the portion of the list in which the word occurred, when portions represent tenths or perhaps eighths.

The second response measure results from asking the subjects to estimate the lag between two items from the list. Lag represents the "distance" between two words as measured by the number of other words that, on the study list, fell between the two test words. For example, if the two words in a test pair occupied positions 15 and 20 of the study list, the true lag is 4. Of course, the subjects are presented many pairs from the list, and for each pair they respond with a number to represent the lag estimate.

The third response measure comes from asking the subject to make recency judgments. Pairs of words are shown on a test (just as if lag judgments were to be requested), and for each pair the subjects designate the word which they believe occurred most recently in the list as presented on the study trial.

A major purpose of Experiment I was to examine the relationships among these three response measures. An acquaintance with two facts is necessary

to understand this purpose. A number of investigators have shown that position judgments have validity, e.g., Toglia and Kimble (1976). On the other hand, three studies (Hintzman & Block, 1973; Hintzman, Summers, & Block, 1975; Underwood, 1977) have demonstrated that there was at best only a slight relationship between true lag and lag estimates for unrelated words. Indeed, in the Underwood study, there was no evidence of substance that the relationship between lag judgments and true lag increased over trials. The question which arises is this: If subjects can make valid position judgments, why cannot the information leading to these judgments be used to make valid lag judgments? More specifically, when two words are presented for a lag judgment, it would seem that the subjects could estimate the position of each, and then take the difference to represent the lag.

Work on temporal coding is a relatively new area of investigation. There is much to be learned about the influence of many variables. The apparent contradiction between the results for lag judgments and those for position judgments evolved from examining the results of separate studies, and these studies differed in a number of ways. Therefore, it seemed necessary to determine if position judgments show validity and lag judgments do not when the tests are made in the same experiment using exactly the same materials and procedures. Experiment I makes this comparison. Actually, all three response measures (recency, lag, position) were used. Furthermore, we asked about changes in performance over trials, and about transfer from one response measure to the other. If common information underlies all three types of judgments, subjects should be able to transfer from the use of one response measure to another with little or no disturbance in the

performance.

All subjects were given five study-test trials. For the groups used to study transfer, the response measure was changed after the first three trials, so that transfer effects will be gauged by the performance on the fourth and fifth trials. The three response measures will be designated by letters: P (position); R (recency); L (lag). Each of the nine groups of subjects may be designated by two letters, the first letter indicating the response measure used on the first three trials, and the second letter indicating the response measure used on trials 4 and 5. The nine groups were PP, RP, LP, RR, PR, LR, LL, RL, PL.

Method

List. A list of 50 five-letter words was used for all conditions. The 50 words were A and AA words from Thorndike and Lorge (1944), and a single random order served as the study order for all conditions and trials. Within the list, 20 pairs of words were identified for use as test pairs for the recency and lag judgments. There were 10 lags (2, 4, 6, 8, 10, 12, 14, 16, 18, 20), with two pairs representing each. Obviously with only two pairs at each lag, we had no intention of examining lag effects in such small steps. Rather, in presenting the results we will speak only of short (2 through 10) and long (12 through 20) lags, each being represented by 10 pairs. The 40 words used to construct the 20 test pairs for lag and recency judgments were also used for the position judgments. Ten words were never tested, although the subjects were not told this. These 10 words occupied positions 1, 2, 16, 21, 22, 27, 33, 42, 49, and 50.

Procedure and subjects. The 50-word list was presented at a 4-second

rate on study trials, using a memory drum. The first test was unpaced for all conditions so that the subjects could understand fully the nature of the tests before pacing was introduced. For the recency judgments, the subjects were given a sheet on which the 20 pairs were listed and they circled one word in each pair to indicate the most recent word. For the lag judgments, the same 20 pairs were presented with a blank after each. The subjects were asked to indicate the number of words which fell between the two on the study trial. They were told that no lag greater than 25 should be recorded. For the position judgments, the subjects were given a sheet of paper on which the 40 words were listed, with a blank after each. The subjects were requested to fill in each blank with a number between 1 and 50 to indicate the position held by the word in the study order. They were further told that they should try to avoid using the same number twice, but this was not monitored by the experimenter. The experimenter monitored each test only to make sure that no omissions occurred.

None of the procedures came as a surprise to the subjects. They were fully informed about the type of unpaced test before the study trial, and again after the study trial, for all trials. All tests after the first were paced at a 6-second rate. Each test word (position judgments) or test pair (recency and lag judgments) was shown for 6 seconds during which the subjects made the decision required. In the few cases where the subject failed to respond within the 6-second interval, the experimenter returned to the pair or word after the completion of the test trial and the subject was required to make a decision. The order of the items on the test differed for each of the first three trials. On trials 4 and 5, the orders used for trials 1 and

2 were repeated. The left-right position of the words in the pairs also varied randomly. When the response measure was changed after the first three trials for the transfer groups, the subjects were completely informed about the nature of the change before the fourth study trial.

A block randomized schedule of the nine conditions was used to assign 20 subjects to each condition. In all experiments to be reported, the subjects were college students.

Results

Position judgments. As a measure of positioning, we used a hit measure. This was defined by a five-position span which included the true position and the two positions on either side of the true position. Thus, if the true position of an item was 25, a hit or correct response was said to have occurred if any one of five numbers (23, 24, 25, 26, 27) was indicated. This was a highly reliable response measure. For example, for the 20 subjects in Condition PP, the correlation between the scores on trial 4 and those on trial 5 was .92.

Figure 1 shows the basic results for positioning. Over the five trials, performance increased from approximately six correct (15%) to over 22 correct (55%) on the fifth trial. The transfer from recency judgments to position judgments (Condition RP), and the transfer from lag judgments to position judgments (Condition LP) was obviously not complete. Still there is evident positive transfer. A comparison of trials 1 and 2 of Condition PP with trials 4 and 5 of Condition LP showed this statistically, F (1,38) = 16.60, $p < .01$. But that transfer was incomplete was shown by the comparison of

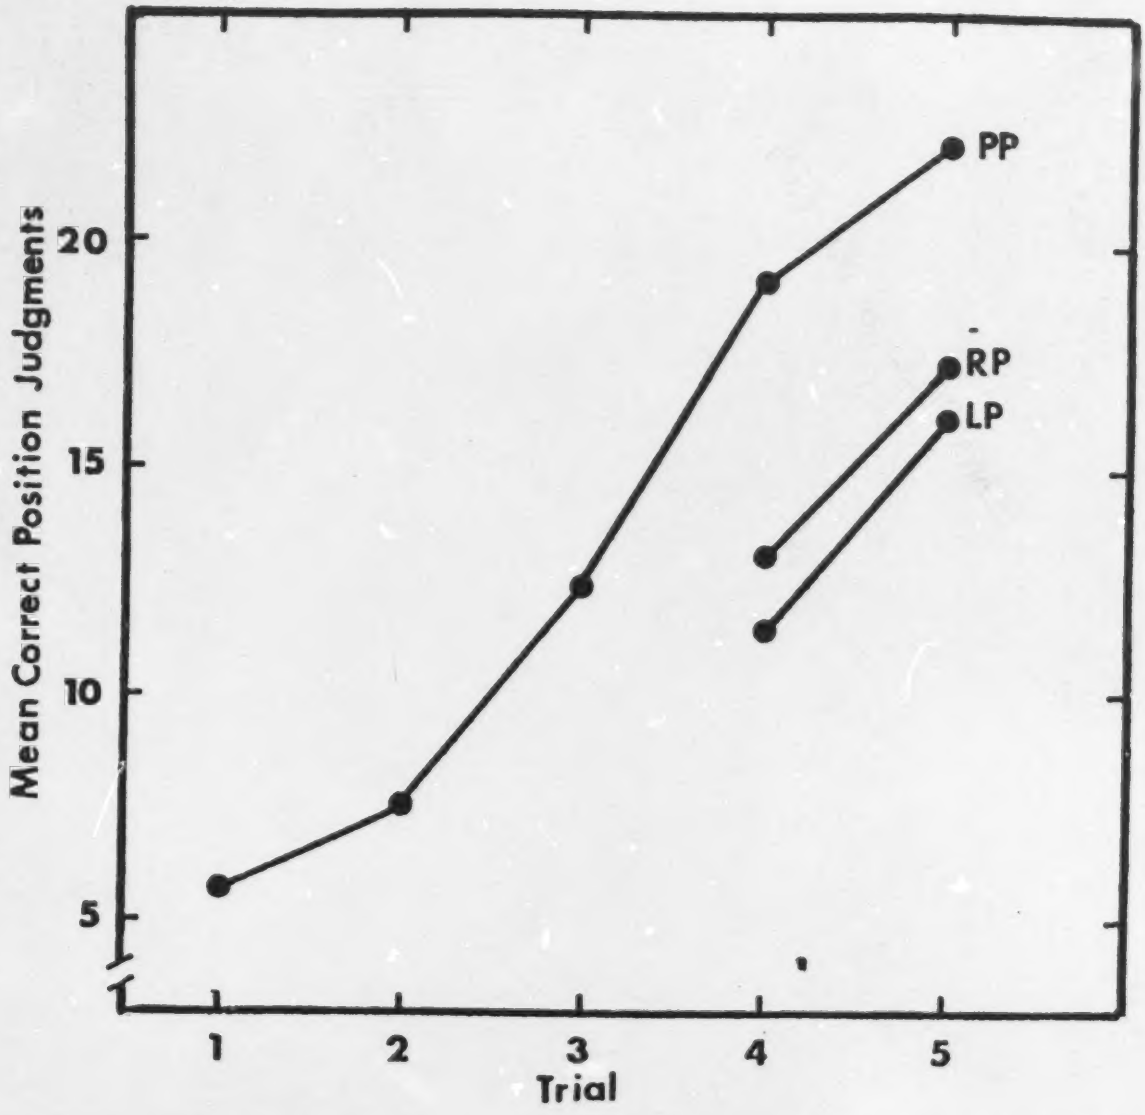

Figure 1. The acquisition of position information over five trials, and the transfer from recency judgments to position judgments (RP), and from lag judgments to position judgments (LP). Experiment I.

Condition PP with Condition RP on trials 4 and 5, F (1,38) = 4.65, $p < .05$. As we will see, the transfer effect is particularly revealing for Condition LP in which the transfer trials were proceded by lag judgments.

We also examined position judgments using the 40 items summed across subjects as the basis for the analysis. A mean position score was determined for each of the 40 words for Condition PP; this was simply the average position assigned each word by the 20 subjects. We then correlated the mean position estimates with the true positions for the 40 words. Across the five trials these correlations were as follows: .66, .74, .89, .92, .95. This obvious increase in the relationship across trials when item scores were used as the entry, supplements the relationship shown in Figure 1 in which subject scores were used as the entry.

Lag judgments. The results for Condition LL are shown in Figure 2. The measure of lag was simply the mean lag judgments, and the values on the ordinate reflect the means per item. The lag length (short and long) appears as a variable in addition to trials. The mean true lag for the short-lag category was 6, that for the long-lag category, 16. On the first three trials there was no consistent difference in the judgments made for long lags and those made for short lags. This confirms the previous findings as described in the introduction to this experiment. The conclusion that lag had no effect on the early trials did not change when the results for all 60 subjects having three lag-judgment trials (LR, LP, LL) were combined.

On trials 4 and 5 of Condition LL, the performance on the short and long lags distinctly separate. An analysis of variance across all five trials showed that lag F (1,19) = 6.44, $p < .05$; trial, F (4,76) = 7.13, $p < .01$; and the lag by trial interaction, F (4,76) = 6.53, $p < .01$, were all reliable. Thus, it seems that some lag learning was exhibited on the fourth and fifth trials. However, there is an asymmetry involved. Judgments of short lags should move lower if learning were occurring. Figure 2 shows that this did not occur. The judgments for long lags did increase as they should if learning occurred. Statistically, then, a lag effect did develop over trials but its nature was not entirely as one would expect.

Figure 3 shows the mean lag judgments on trials 4 and 5 combined for the three relevant conditions. As a simplifying step, trials have been omitted from the plot. The trial effect was reliable statistically, but there was no interaction involving trials. The fact that conditions differ is not very meaningful, since it tells us only that mean judgments differ as to level of responding. The central concern are the differences in discrimination between short and long lags as a function of conditions. The lag effect was significant statistically, F (1,57) = 30.04, $p < .01$, but the lag by conditions interaction was far from reliable, F (2,57) = 1.51, $p > .05$. This means that the transfer was complete for both conditions; the discrimination between long and short lags was as great for the two transfer conditions (RL and PL) as it was for Condition LL. At the same time, we must point out that the discrimination, while statistically reliable,

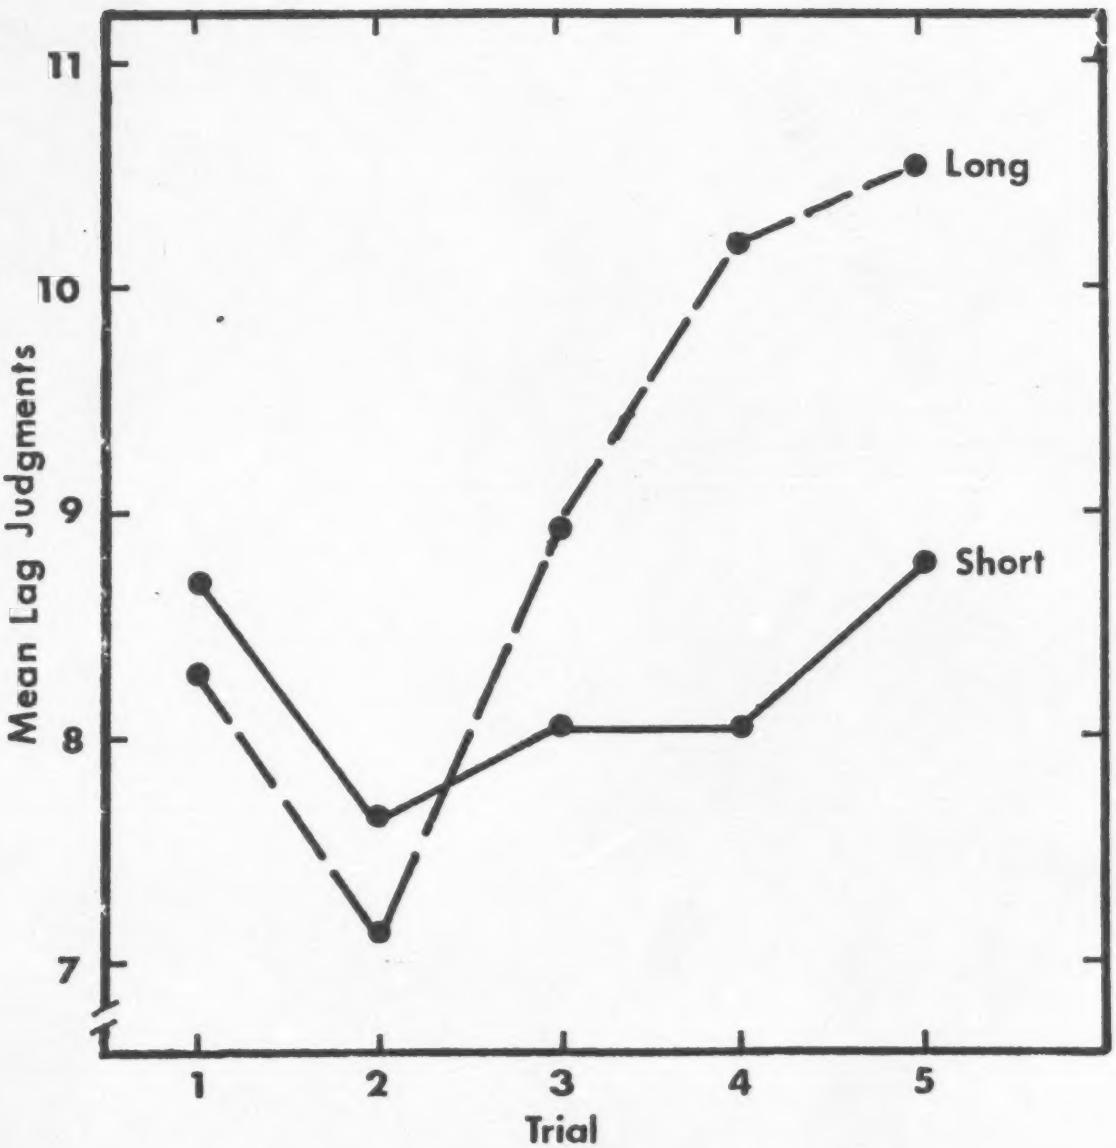

Figure 2. Lag judgments as a function of lag (long and short) and trials. Experiment I.

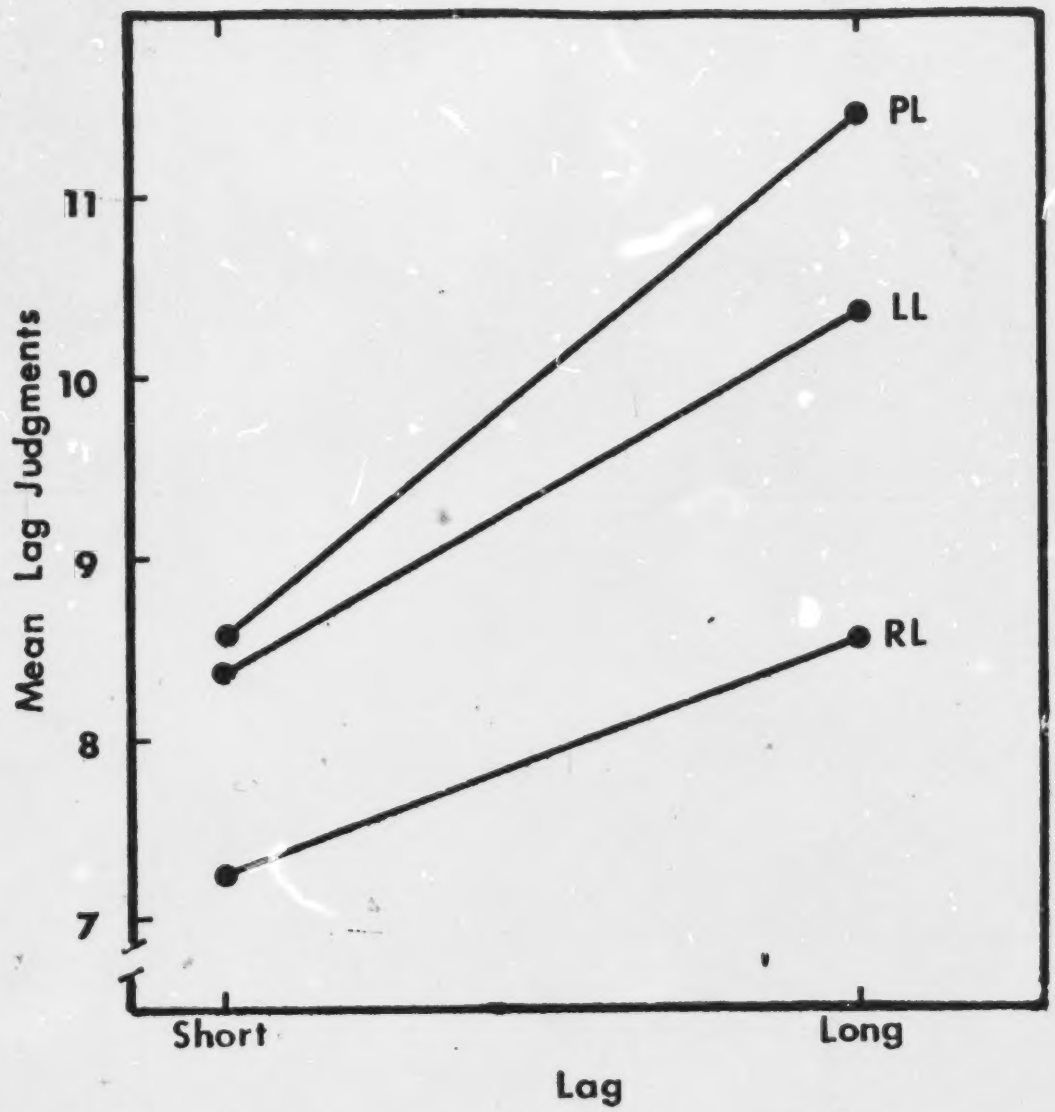

Figure 3. Transfer from position judgments to lag judgments (PL), and from recency judgments to lag judgments (RL) as a function of lag. Experiment I.

is, in an absolute sense, not very impressive.

Recency judgments. The recency measure was simply the mean number of correct recency judgments. These are plotted in the upper section of Figure 4 for Condition RR as a function of lag length and trial. Scores could range between 0 and 10, with a score of 5 representing chance responding. On all five trials there was a clear separation as a function of lag F (1,19) = 13.97, $p < .01$. Performance improved over trials, but the F for the interaction between trials and lag was less than one. The fact that lag influenced recency judgments is a contradiction of previous work (Underwood, 1977, Experiment 11) where lags of 3, 6, 9, and 12 did not influence performance reliably over three trials. For reasons which we do not understand, our previous result was anomolous with regard to the effect of lag differences on recency judgments. Every study we report here in which lag was varied showed an effect of lag on recency judgments. The effect is often not large, but it is always present. It occurred on the first three trials of Conditions RL and RP, although we have not shown these data here. Figure 4 shows that the effect does not increase in magnitude over trials. On the fifth trial the values, when translated into percentages, show about 76% correct recency judgments for the short lags, 83% for the long.

The lower panel of Figure 4 provides the transfer data. The line for Condition RR simply represents the mean of the two lines in the upper panel. Although lag was not included in the plot, it can be said that there was a reliable lag effect on trials 4 and 5 for Conditions PR and LR. Analysis

of the data for the three conditions for trials 4 and 5 showed that only trials was a reliable source of variance. Neither conditions nor the conditions by trials interaction approached significance, the F being less than one in each case. We must conclude that the transfer from position judgments to recency judgments, and from lag judgments to recency judgments, was essentially perfect or complete. That the performance for Conditions PR and LR was a little lower than that for Condition RR may result from a performance disturbance, i.e., adjusting to the paced responding using a new response indicator. That transfer was nearly perfect is of some moment in view of the fact that on the first three trials under Condition LR the subjects gave little evidence that they were distinguishing between long and short lags when they made their lag judgments. Obviously they were acquiring information which allowed them to make correct recency judgments as readily as did those subjects who had been making recency judgments on all trials.

Correlations. Correlations may be used to supplement the transfer evidence with regard to the commonality of the information underlying the three different response measures. The correlations in Table 1 include those within tasks as well as those between tasks, where task refers to the different response measures. Scores for various trial combinations were used, namely, trials 1 through 3 combined, trials 4 and 5 separately, and trials 4 and 5 combined. For recency and position judgments, number of correct responses was used to index performance. Lag judgments do not give directly a measure of the goodness of performance. To provide such a measure,

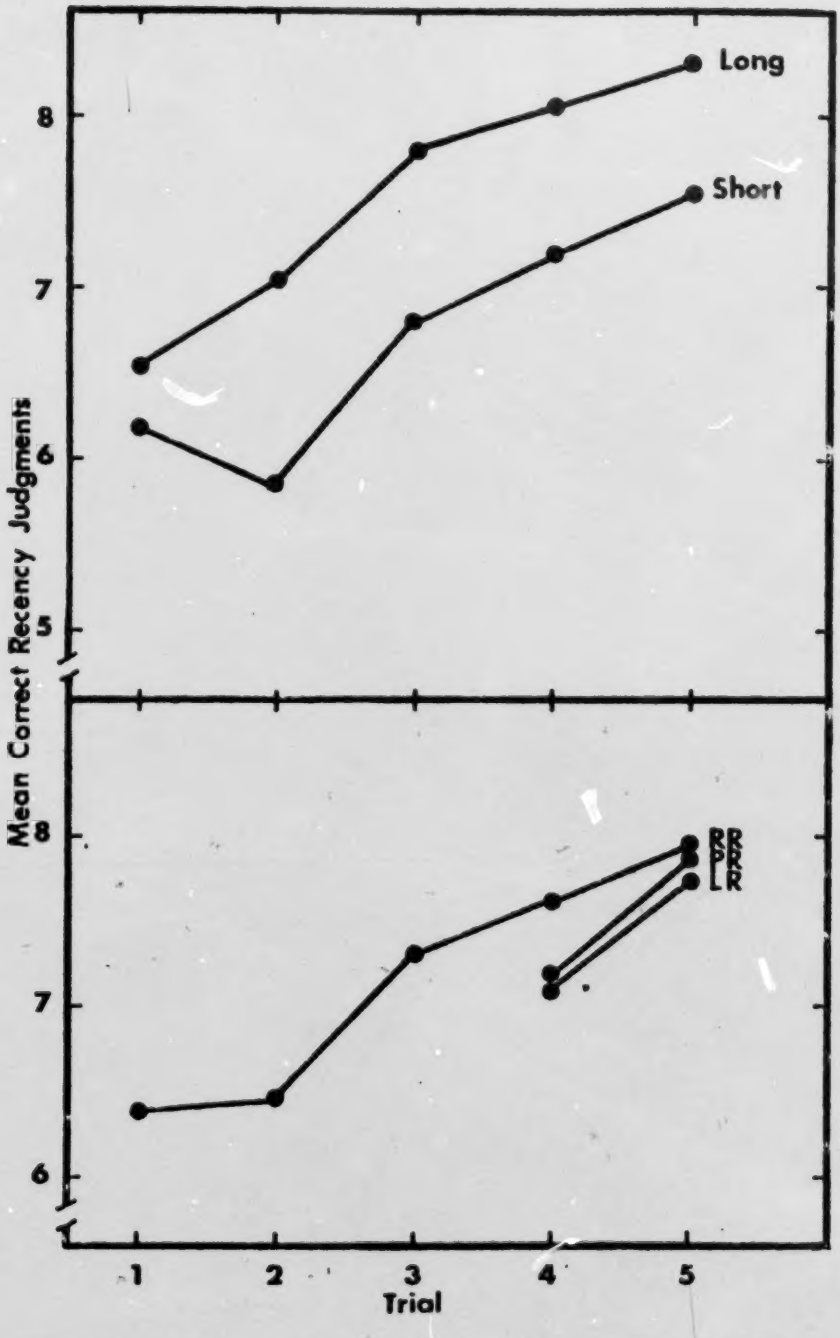

Figure 4. Recency judgments as a function of trials and lag (upper panel), and transfer performance (lower panel) from position judgments to recency judgments (PR), and from lag judgments to recency judgments (LR). Experiment I.

Table 1

Intratask and Intertask Correlations on Various Trial Combinations

(Decimal Points Omitted)

Condition	(1-3) x (4)	(1-3) x (5)	(1-3) x (4-5)	(4) x (5)
LL	12	35	30	70
RL	22	60	51	16
PL	83	61	84	47
PP	74	59	67	92
LP	43	41	45	71
RP	27	2	14	54
RR	74	54	76	71
PR	46	51	59	51
LR	53	41	48	70

we used the sum of the judgments for the 10 long-lag pairs minus the sum of the judgments for the 10 short-lag pairs. This is a slope measure in which the greater the slope (difference between short- and long-lag judgments) the greater the sensitivity to lag differences.

The correlations for Condition LL, PP, and RR give evidence on the reliability of the response measures. Under Condition LL, the performance scores for the early trials showed a very low relationship with the scores on the late trials. This corresponds to the lack of a lag effect on the early trials for this group (Figure 2). The reliabilities for Conditions PP and RR were quite high. For all correlations a value of .45 is required for the .05 level of significance.

If the intertask correlations are to allow for a meaningful interpretation, they should be low (as will be explained shortly). Furthermore, the most critical correlations would be those between trials 1-3 combined and trial 4. This is not a completely uncontaminated relationship because the subjects were told before the fourth study trial that the test would change. Knowing this, the subjects might have changed their encoding habits on the fourth study trial. High correlations between the scores on trials 1-3 and those on trial 4 do not yield a clear interpretation. A high correlation could mean that the underlying information used is the same for two response measures, but it could also mean that the information differs in kind but there is a correlation in terms of the rate of acquiring the two types of information. Therefore, only low correlations can be used in an analytical way.

With the above in mind, we note first that there seems to be little

commonality between recency judgments followed by lag judgments (RL), when the correlation is based on trial 4. However, the subjects apparently find some way to use the information acquired over the first three trials because the correlation between trials 1-3 and trial 5 is quite high. Under the reverse order of tasks (Condition LR), a relationship is apparent throughout, although the magnitude of the correlations is certainly not high. The only other condition to which we will call attention is Condition RP. The scores on these two tasks (recency, positioning) do not correlate substantially at any point. However, under the reverse order, Condition PR, the relationships were all reliably greater than zero. Thus, it would appear that subjects having position judgments and lag judgments before recency judgments can make use of at least some of the information for making recency judgments, but this information is of a different type from that acquired when recency judgments are learned initially. In short, there is ambiguity resulting from different findings associated with the different transfer orders. As a consequence, the correlations in Table 1 have not provided us with the critical data that in the abstract it seemed they should.

Discussion

The experiment confirmed the inference that one group of subjects may make very valid position judgments of a list of items presented singly, while another group has great difficulty in making valid lag judgments after studying the same list. Valid lag judgments begin to emerge only after several study and test trials. We have failed to confirm a previous finding that lag was irrelevant for recency judgments; our present data show clearly that lag length is positively related to correct recency judgments.

The data showed that there was rather heavy transfer from the use of one response measure to the use of another. The transfer was nearly perfect in two of the three cases (Figures 3 and 4), and roughly 50% in the third (Figure 1). In the latter case the transfer was from recency judgments to position judgments and from lag judgments to position judgments. The transfer in the reverse directions was essentially 100%. We will not speculate about the reason for the asymmetry. Rather, we wish to consider briefly the interpretation that might be given to the rather heavy transfer from one response measure to the other. We prefer to interpret the results to mean that the information involved in making decisions overlaps appreciably for the three response measures. The subjects draw on the same types of information regardless of the response measure. There are alternative interpretations, one of which will be described.

We might assume that during the learning trials the subjects acquire a number of different types of information (each independent of the other) about the words in the list, and this information is acquired regardless of the response measure being used. We would further assume that the amount of information of each type which is accumulated is correlated. We might name the three types of information P, R, and L, each of which is appropriate to its own response measure only. When subjects are transferred from one response measure to another, they simply change the type of information selected from memory, choosing the type needed to fit the demands of the new response measure. The consequence would be high positive transfer from one response measure to another, and positive correlations among the scores on

the various tasks. This is a rather complicated interpretation but we know of no evidence that would deny it. If we favor the earlier interpretation, it would be due to the simplicity of the interpretation. Neither interpretation adequately accounts for some of the details, e.g., the asymmetry.

Experiment II

Despite the fact that subjects did show some learning in Condition LL in Experiment I, the fact remains that they found it a very difficult matter to discriminate the "distance" between two items from the list. We do not know whether subjects tried to learn this distance directly, or to learn by indirect means, e.g., learning the position of each item and then taking the difference as a lag estimate. In any event, it seemed worthwhile to try to devise a situation in which accurate lag judgments should be easily acquired on the first study trial. We presented the subjects with a 30-word list in which the first six words were all names of metals, the second six all names of animals, and so on. In short, there were 5 concepts used, each with six instances, blocked in the presentation list. It seemed to us that under these circumstances even a poor subject would assign short lag values if two words from the same concept were tested. To try to produce valid lag judgments quickly was not the only matter of interest in the study; we turn to these other matters.

In the above situation, it is possible that subjects would be quite able to make valid lag judgments for two items from within a concept and have difficulty in making recency judgments for the same two items. Acquisition of the particular order of the instances within the category could be difficult, hence, recency judgments could be difficult. In the experiment, sub-

jects were also asked for lag and recency judgments for pairs of items in which the two words in the test pairs represented different categories. The lag for four of these pairs was two, in which case it was necessary that one of the words be the fifth or sixth item in one category, and the other be the first or second item in the following category. The tests made for items within categories also always had a lag of two. So, there were within-concept or within-category tests with lags of two, and between-concept tests with lags of two. We also included between-concept tests with long lags (varying between 8 and 23).

What did we expect from the recency and lag judgments for test pairs in which each of two concepts is represented? We believed that serial learning would play an important role. If a subject learned the order in which the six concepts was presented, correct recency judgments should be easily "deduced", and for the short lags, the lag judgments should become valid rapidly. In fact, we felt that the lag judgments on the short-between tests could be more accurate than those for the within-category tests. The unknown in this matter concerns the rapidity with which the subject learns the serial order of the words within categories.

Although the above comparisons are of some interest, we will be most concerned with the comparison of temporal codes developed for the blocked category list and those developed for the same items when they are randomly assigned to positions within the list. Also, we will be asking about transfer from lag judgments to recency judgments and for the reverse order. Position judgments were not included. Thus, with blocked (B) ordering of the category instances within a list, and the unblocked (UB) or random ordering

of the words, eight groups were needed: B-LL, B-LR, B-RL, B-RR, UB-LL, UB-LR, UB-RL, and UB-RR.

Method

Lists. The general nature of the B list was described above. There were 30 words consisting of five categories of six instances each, taken from the tables provided by Battig and Montague (1969). The order of the categories was metals, animals, cloths, sports, and musical instruments. In addition, there were two primacy items and two recency items. After the words were ordered within the list, we chose four pairs to be used for within-category tests. These consisted of the items holding positions 3 and 6 in the metals category, 2 and 5 in the animal category, 3 and 6 in the sports category, and 1 and 4 in the instruments category. Four between-category pairs were chosen as the between-category, short-lag tests (between-short). The lag was always two. From the remaining 14 words, we constructed six test pairs to be used for the between-category tests with long lags (between-long). The six lags were 8, 9, 13, 13, 16, and 23, with a mean of 13.7. Two words (in addition to the primacy and recency words) were not used for the tests; one of these words was the first word in the metals category (third word in the list), and the other was the last word in the cloth category.

All of the above description refers to the blocked (B) list. We found it possible to order the items in an unblocked (UB) list so that the 14 test pairs used were exactly the same pairs as those used for tests of the blocked lists. Furthermore, the lags for the pairs in the two lists were essentially the same, pair by pair. All short lags were two for both lists. Three of the long lags were identical for both lists, but the remaining three varied

by one (8 vs. 9; 22 vs. 23; 16 vs. 17). In only one case did two items representing the same category occur in adjacent positions in the unblocked list.

Procedure and subjects. The procedure was much the same as for Experiment I except that all test trials were paced (6-second rate). There were five study-test cycles for all subjects, and for the transfer conditions, three study trials were given on the first-used response measure, two on the second. The subject was always completely informed. Each of the eight conditions was represented by 24 subjects assigned to conditions by a blocked-randomized schedule.

Results

Lag judgments. It will be remembered that the lag for within-category tests was always two, and this was also true for the between-short tests. However, since there were only four pairs in each class, and since the words were not rotated across the classes, interclass comparisons are questionable. The results will emphasize comparisons between B and UB conditions for each class because the judgments were made on exactly the same pairs for such comparison.

The mean lag judgments for within-concept tests, and those for the between-short tests, are shown in Figure 5. As may be seen in the left panel, the within judgments are very accurate for Condition B-LL. The means are only slightly higher than the true lag (two), and the level of

responding is far more accurate than the level for the subjects under Condition UB-LL. Although the trend across the five trials for Condition UB-LL indicates a lowering of the estimates, the means on the five trials did not differ statistically ($F = 1.09$). Thus, as in Experiment I, lag judgments did not move toward the true length of short lags.

Transfer from recency learning (first three trials) to lag judgments (trials 4 and 5) appears to have been essentially complete for Condition B-RL. This is to say that the evidence which allowed the subjects in Condition B-LL to make very accurate lag judgments on the within pairs was also assimilated by the subjects in Condition B-RL. In fact, as we will see later, correct recency judgments developed more slowly than did lag judgments. On the transfer tests, the subjects in condition UB-RL performed at about the same level as did the subjects in Condition UB-LL, although it was not statistically lower than the level of responding on trials 1 and 2 for Condition UB-LL, $F(1,46) = 3.71$, $p > .05$. This would seem to indicate there was no transfer from recency judgments to lag judgments. On the surface this conclusion appears in conflict with that of Experiment I where it was said that transfer from either recency judgments or position judgments to lag judgments was complete (as gauged by the performance on Condition LL). However, that decision was based on examination of a slope measure derived from the scores on short and long lags. Were we to follow that procedure here, we would reach the same conclusion as was reached in Experiment I because judgments on long lags do increase across trials (see later) and the level of the transfer performance is equivalent to that shown under Condition UB-LL.

The right panel of Figure 5 shows the results for the pairs with lags

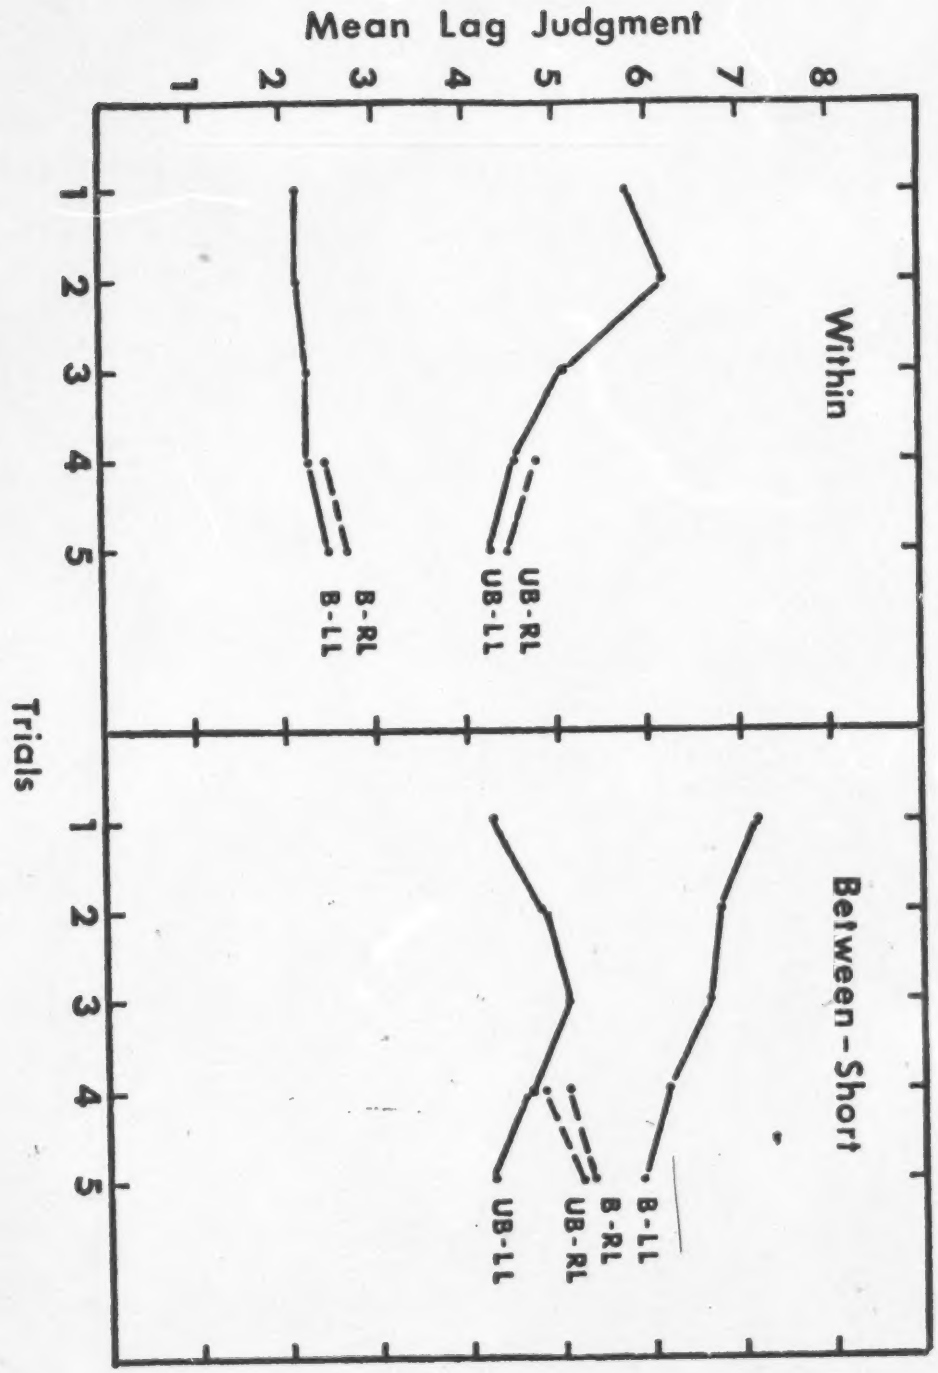

Figure 5. Judgments of lag when the lag was always two for items within a concept, and for items from two different concepts. The dotted lines represent the transfer performance. See the text for further elaboration. Experiment II.

of two in which the pair members were items from different categories (between-short). There was very little change over trials for Conditions B-LL and UB-LL ($F = 1.46$). Essentially, the subjects in Condition UB-LL performed at the same level as they did on the within pairs. The big change occurred for Condition B-LL in that the lag judgments were higher than those given by the subjects in Condition UB-LL, $F\ (1,46) = 9.35$, $p < .01$. Neither trials, nor the interaction between trials and conditions, was reliable statistically. Because no learning occurred in either group (as would be indicated if the means decreased reliably over trials), the transfer data must be viewed as representing no transfer.

We have not graphed the results for the items classed as between-long. The data for them showed that the mean judgments were higher for Condition B-LL than for Condition UB-LL, $F\ (1,46) = 6.03$, $p < .02$. The five means corresponding to trials 1-5 for Condition B-LL were 8.44, 8.87, 10.06, 9.49, and 9.46. For Condition UB-LL, the corresponding values were 5.84, 6.58, 6.98, 7.64, and 7.78. The mean judgments (both conditions combined) increased over trials, $F\ (4,184) = 5.09$, $p < .01$, but the F for the interaction between conditions and trials was less than one.

There was no transfer from recency judgments to lag judgments under Condition B-RL for the long-lag items. The sum of the judgments on the two transfer trials was actually slightly lower than the sum for trials 1 and 2 of Condition B-LL. On the other hand, the transfer appeared to be complete in Condition UB-RL. However, as will be discussed later, we are concerned about the validity of these comparisons.

Recency judgments. The mean correct recency judgments for the two

short-lag conditions are shown in the two panels of Figure 6. Because there were four possible correct responses, a mean of two would represent chance responding. The results in the left panel show that the mean correct recency judgments across five trials for two items within a category of six instances (Condition B-RR) were greater than the mean for the same items in a randomized list (Condition UB-RR), F (1,46) = 5.96, $p < .02$. The difference between these same two conditions for the between-short tests (right panel) was even more marked.

Transfer was statistically complete in three out of the four cases. The only exception was under Condition UB-LR for the within tests. Comparing the performance on the two transfer trials of this condition with the performance on trials 1 and 2 of Condition UB-RR showed that there was positive transfer, F (1,46) = 7.95, $p < .01$. But, comparing the performance on the transfer trials with trials 4 and 5 of Condition UB-RR showed that the transfer was not complete, F (1,46) = 6.23, $p < .05$. The general implication of these transfer findings is that although subjects learned very little about lag lengths when making lag judgments, they did acquire information which allowed them to make recency judgments at a level that approached that expected if recency judgments had been practiced directly.

The results for the between-long test pairs were very similar to those seen in the left panel of Figure 6. Performance under Condition B-RR exceeded that under Condition UB-RR, F (1,46) = 9.97, $p < .01$. Responding

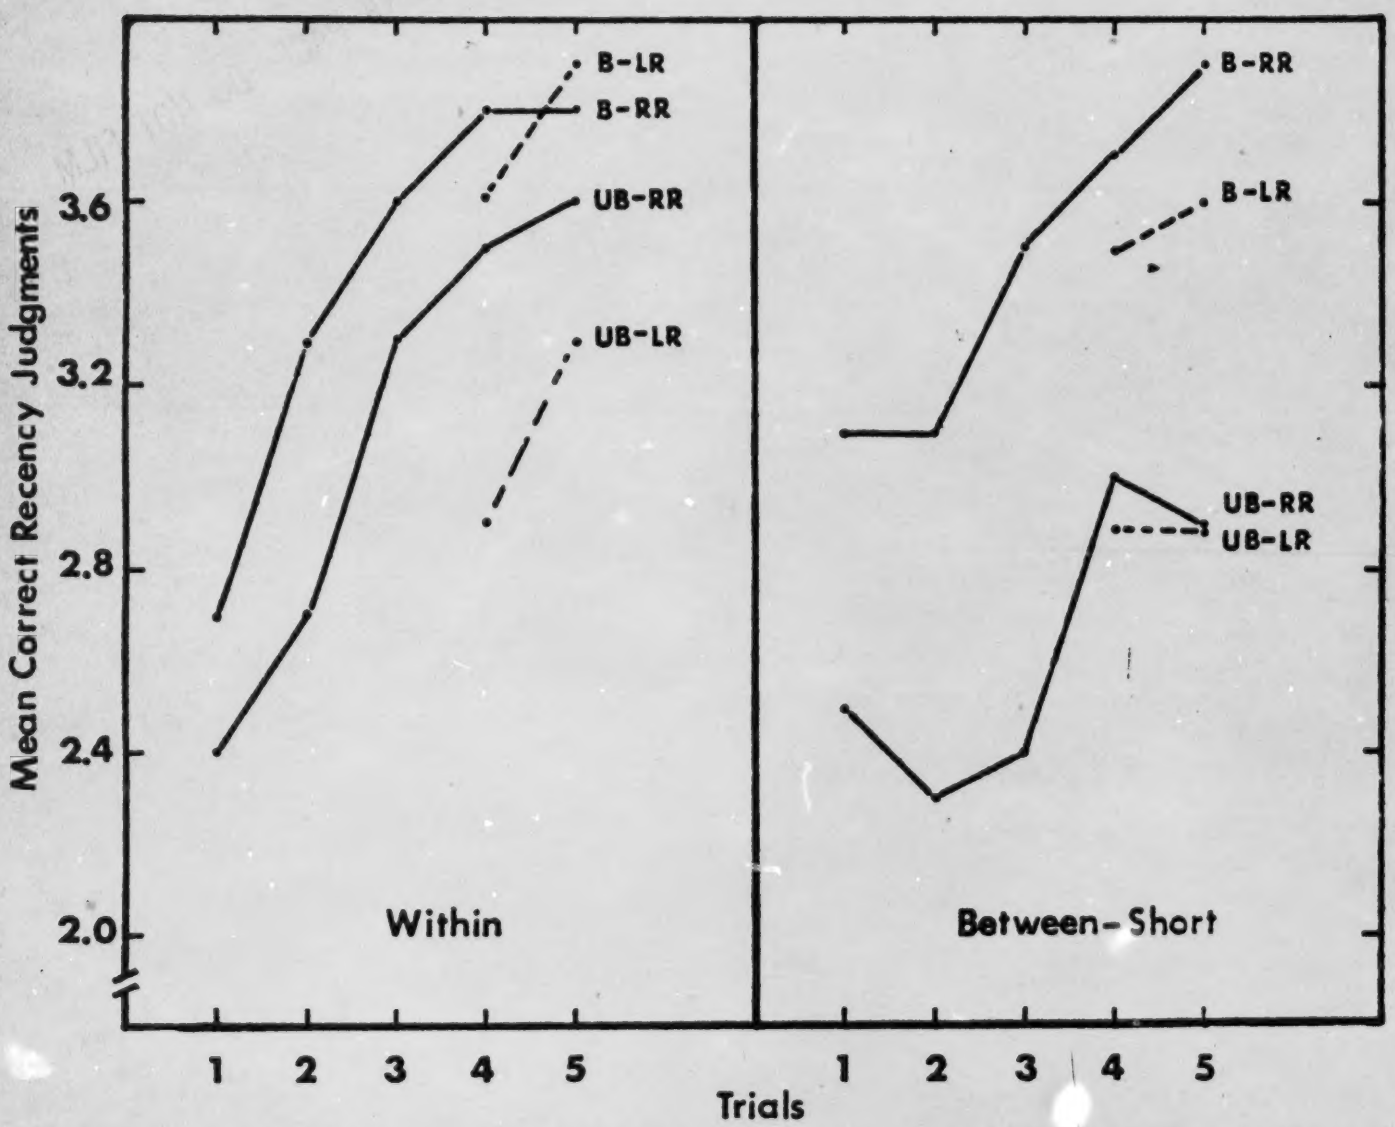

Figure 6. Recency judgments for the same items used to produce the lag judgments in Figure 5. Experiment II.

was nearly perfect on trials 4 and 5 under Condition B-RR. Transfer for Condition B-LR was complete, the performance for the subjects in this group also being nearly perfect on trials 4 and 5. Transfer under Condition UB-LR was complete statistically in that performance on the two transfer trials (83.3% correct recency judgments) was not reliably less than the performance on trials 4 and 5 under Condition UB-RR (89.9%).

Discussion

The results have shown that lag judgments can be highly accurate when certain kinds of information are available to the subject. In particular, if the two test words are from a block of six words representing the same category, lag judgments are very accurate. The subjects apparently learn the nature of the blocked list structure on the first study trial, and from this is can be "deduced" that if two items in a test pair are both instances of the same category, the lag must be short. This may seem to be a trivial finding, but we have become somewhat skittish about attributing logical capacities to the subjects when lag judgments are involved. For example, it seemed obvious that if a subject would make valid position judgments, accurate lag judgments could be deduced from the position information. This turned out not to be true. In any case, for the present experiment, short lag judgments were brought under control by using the blocked category lists and testing pairs of items from within a block.

The other two classes of items produced results in which there are a number of puzzles. The mean lag judgments for the subjects in Condition B-LL for between-short pairs were higher than those for the same pairs appearing in Condition UB-LL (Figure 5, right panel). Although neither

group improved significantly across trials, we might choose to conclude that because the level of responding was lower (hence, closer to true lag) for the subjects in Condition UB-LL than for those in Condition B-LL, performance was better or more accurate. We might, further, talk about differences in interference which results from conceptual associations in the two conditions. There are problems which follow if it is concluded that the performance of the subjects in Condition UB-LL was superior to that of the subjects in Condition B-LL. If this was true, why didn't the subjects in Group UB-LL get more and more accurate with practice? Or, if this was true, why didn't the same relationship occur for the long lags? The judgments for long lag pairs for the subjects in Condition B-LL were higher (more accurate?) than were those for the subjects in Condition UB-LL. We are inclined to believe that such conclusions cannot be drawn with confidence, and that in some way the conditions produce differences in the level of responding although there are no differences in lag knowledge.

It is probably correct to presume that the subjects in Condition B-LL knew they were accurate for the within-category tests, but that they lacked confidence in their ideas of the absolute lags for the other test pairs. Given this situation, the question is why did the subjects in Condition B-LL choose a higher responding level for the between-short pairs than did the subjects in Condition UB-LL, the latter subjects having no class of lag judgments in which they could feel confident. It was almost as if the subjects in Condition B-LL said to themselves: "I know my short-lag decisions for the pairs of items from the same category are about perfect; since I don't seem to have such feelings for the other pairs, they must have long

lags." As a consequence of such thoughts or similar ones, the subjects may have assigned values to lag judgments that were higher than those which were assigned by the subjects in Condition UB-LL. Because of our concern that such contrast-like effects may have been involved in the judgments, we are uneasy about the transfer effects from recency judgments to lag judgments for Condition B-RL, when the comparison involves Condition B-LL. The subjects in Condition B-RL may not have been influenced by the contrast effects, so that the level of responding would not be expected to be comparable to that for Condition B-LL even if lag knowledge was the same. Hence, to determine if there was or was not positive transfer by comparing performance on the first two trials of Condition B-LL with trials 4 and 5 of Condition B-RL may lead to erroneous conclusions about transfer. On the other hand, we believe the transfer results for Condition UB-LL and UB-RL would not have this problem, although, except for long lags, there was no evidence that judgments improved over trials, hence there was no opportunity to assess transfer effects.

In the introduction we pointed out that if subjects in Condition B-LL mastered a five item serial list (metals, animals, cloths, sports, instruments), lag judgments should improve rapidly over trials for the tests involving items from different categories. The subjects could attach the numbers 1 throught 5 to the five categories in order. Given this learning, relative lag judgments should have been perfect if the subjects, on the tests, identified the numbers of the two categories represented and then took the difference. Of course, this is behavior we had originally bequeathed to subjects as a means of making lag judgments from position knowledge, and we

found that a subject apparently cannot or does not go through such mental manipulations. The subject can make accurate position judgments without being able to make accurate lag judgments.

Transfer from lag judgments to recency judgments was very high, just as we had found in Experiment I. This transfer was heavy even if the subjects showed no increase in the accuracy of lag judgments over the three "training" trials. The subjects, therefore, were acquiring most of the information needed to make recency judgments while they were trying without much success to make accurate lag judgments.

Recency judgments were consistently better for the subjects in Condition B-RR than for those in Condition UB-RR. The differences were particularly large for the between short tests (right panel, Figure 6). Why should these two conditions differ? One strong possibility is that the serial learning of the five categories could be used to implement between-category recency judgments. If the successive categories are numbered 1 through 5 when the list is blocked, the most recent instance is given directly by the two numbers associated with the two words in a test pair. The responding in this case does not require that the subject take a difference between values as seems to be true for lag judgments. This explanation, of course, will not handle the difference for the two groups for the within-category tests (left panel, Figure 6). Serial learning within the four concepts may be involved in the latter case, but performance seems to improve more rapidly over trials than would be expected if the serial learning of four lists of six items was involved. It may also be possible that the cause of the difference lies not in a facilitation under Condition B-RR, but in inhibition

under Condition UB-RR. The inhibition might result from greater associative interference for the random list than for the blocked list, but just how this interference operates on recency information is not known.

We have chosen not to pursue further in this paper the interpretative problems which have arisen by our attempts to understand the relationships among different response measures used to index temporal coding of memories. Other problems have engaged our attention. In the experiments on these problems, we have always used recency judgments as the response measure.

Experiment III

In an experiment reported earlier (Underwood, 1977, Experiment 5), it was shown that exposure duration of items did not influence the accuracy of position judgments. If this finding has generality, it would most assuredly influence the nature of theories about temporal coding. We believe that associative learning is primarily responsible for the learning of correct recency judgments and correct position judgments, and associative learning is normally quite sensitive to rate or exposure-time manipulations. The shortest exposure duration used in our earlier study was 5 seconds. It is possible that the major changes in temporal coding associated with changes in exposure duration occur with durations shorter than 5 seconds. We felt it necessary to determine the role of relatively short exposure durations of the items on recency judgments.

The usual method of manipulating exposure duration confounds study time and retention interval. This is true even if tests are given immediately after the last item is presented for study The length of the retention interval for items in the initial part of the list is directly related to

exposure duration. There is another problem of confounding which arises when recency judgments are used as the response measure. Although the lag (number of items) between two test items is constant regardless of exposure duration, the time between the two items varies as exposure duration varies. We were concerned about these problems of method, and we determined to decide if they were matters of importance for temporal coding. Therefore, we used one condition as representative of the traditional approach, and another in which the confoundings were removed. The exposure durations were of three lengths; 1, 2, and 3 seconds.

Method

Lists. Three lists of 36 words were used. From a pool of 108 A and AA, five-letter words, we constructed three lists such that words were assigned to lists and to positions within the lists randomly. The three lists were arbitrarily designated 1, 2, and 3, and these numbers represent the exposure durations given them. Eight pairs from each list were designated as short-lag test pairs (lags of 2 or 3), and eight were designated as long-lag test pairs (lags of 10, 11, or 12). Items 1, 2, 35, and 36 in each list were never tested. It should be clear that the test pairs came from exactly the same positions in all three lists. Thus, short-lag test pairs were made from the two words occupying positions 5 and 9 in the three study lists, the test pairs being party-greet, awake-count, and offer-happy for the three lists in order.

The usual method of manipulating study time or exposure duration will be called Condition V to indicate that the retention interval and the "crowding" varied as the exposure duration varied. In the condition used to keep these two factors constant, Condition C, we set the memory drum to

provide a 1-second exposure under all conditions. The first nine exposures for each list as given in Table 2 illustrate the method used to remove the confoundings. It can be seen that items are exposed for 1, 2, and 3 seconds, and that the total time from beginning to the end of each list was the same for all conditions. We are, of course, assuming that the arithmetic problems will prevent rehearsal, and at the same time will not interfere with the learning required to produce correct recency judgments.

Procedure and subjects. The subjects were fully informed about the recency judgments they were to make after each study trial. All tests were unpaced. The subject was given a test sheet on which the 16 pairs were listed; the requirement was to circle the most recent word in each pair, guessing if necessary. During study trials involving the arithmetic problems, the subjects solved each equation, speaking the answer aloud within the 1-second period. Three study-test cycles were used for all lists.

One group of 24 subjects was given the three lists under Condition V, and another group of 24 subjects was given the three lists under Condition C. Within each group, three orders of the three lists were used, 1-2-3, 2-3-1, and 3-1-2. Eight subjects were assigned to each order within each condition. Assignment of subjects to condition and to list order followed a block-randomized schedule of the conditions.

Results

Recency judgments. The recency judgments were always reduced to reflect a base of eight (the number at each lag). Hence, a mean of 4 would

indicate chance responding. We may first observe the upper panel in Figure 7, where exposure duration is related to condition. Although the performance under Condition V is better than under Condition C, the difference is far from being reliable, F (1,46) = 1.20, $p > .05$. Differences in exposure duration were associated with differences in the number of correct recency judgments (F = 17.12), but the interaction between exposure duration and conditions was unreliable, F (2,92) = 1.56, $p > .05$. It appears that our concern about differences in the retention interval as a function of exposure duration was a needless one.

The lower portion of Figure 7 relates exposure duration to lag. The differences as a function of lag are quite apparent (F = 33.13). Neither of the plots shows trials as a variable, but as usual, performance increased over trials; the mean values, summed across the other variables were 4.50, 5.08, and 5.36 for the three trials in order. None of the interactions among any of the variables approached statistical reliability. Finally, we could find no effect of the orders in which the three lists were learned.

Correlations. The evidence as described above seemed quite clear in demonstrating that Conditions V and C did not produce differences of consequence. Differences in the length of the retention interval inherent in the method used in Condition V had little influence on performance. Similarly, it appears that requiring the subject to make simple arithmetic calculations under Condition C had little influence. In a further effort to detect differences between the two conditions, we calculated all possible

Table 2

The First Nine Exposures (1-second rate) for Each List of Condition C, Experiment III

Exposure	List 1	List 2	List 3
1	rough	flash	great
2	4 + 2 = ?	flash	great
3	8 - 6 = ?	1 + 7 = ?	great
4	fancy	smart	steal
5	2 + 7 = ?	smart	steal
6	7 - 5 = ?	4 + 5 = ?	steal
7	claim	grown	event
8	5 - 1 = ?	grown	event
9	3 + 4 = ?	5 - 3 = ?	event
	etc.	etc.	etc.

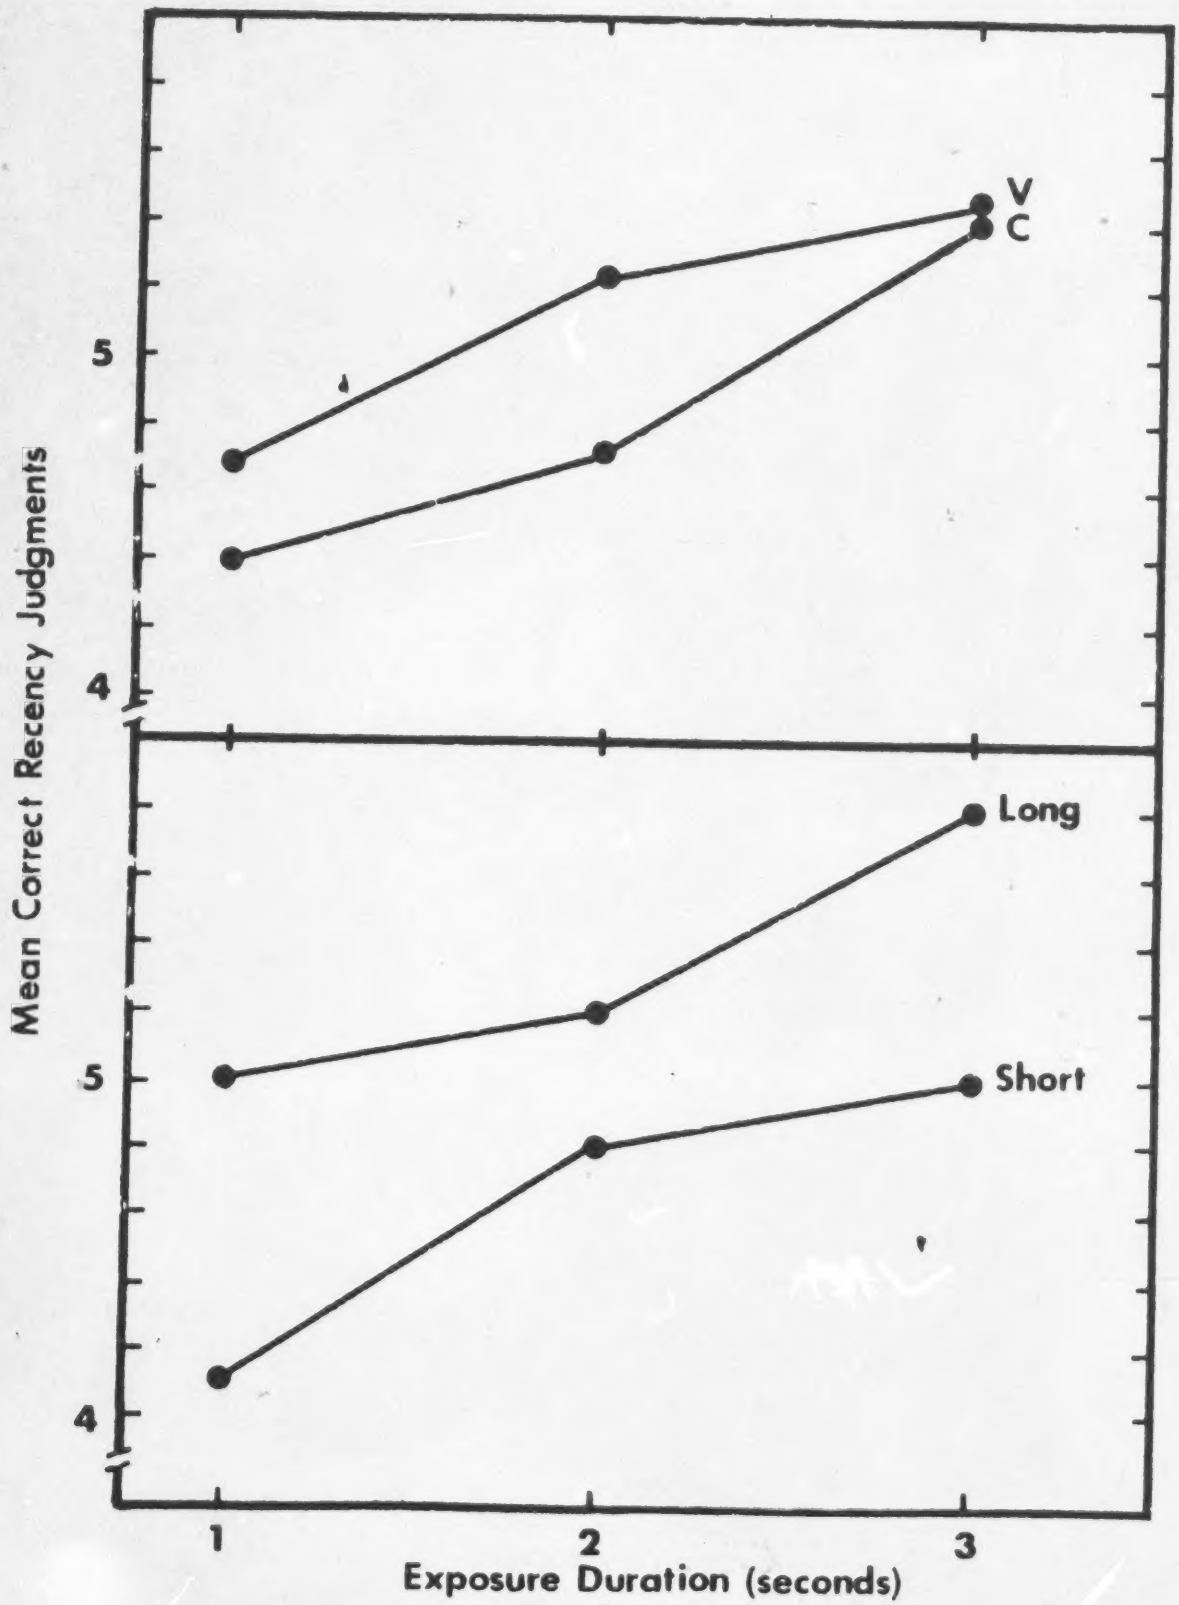

Figure 7. Recency judgments as a function of exposure duration and variable (V) or constant (C) conditions (upper panel), and the relationship with lag (lower panel). Experiment III.

intracondition correlations, where lag and list were kept separate, thus allowing 15 different correlations. Of these 15 correlations for Condition V, 11 were reliably different from zero ($p < .05$, 22 df, $r = .40$). For Condition C, only two of the 15 were reliably different from zero. When the 15 correlations for each group were aligned in two columns so that list length and lag were matched, the correlations were higher in every case under Condition V. The mean r was .19 for Condition C, and .46 for Condition V.

Our post hoc interpretation of these differences points to different performance requirements for the two conditions. The performance requirements were more varied under Condition C than under Condition V. With the 1-second exposure duration for Condition C, the subjects solved two arithmetic problems between each word; with the 3-second exposure duration they solved none. There might, therefore, be an interaction between subjects and recency learning associated with the differences in the performance requirements, and this is reflected in the correlations.

Discussion

With respect to the central purpose of the study, the results were unambiguous. Correct recency judgments increased directly with increases in exposure duration from 1 through 3 seconds. The earlier finding that position judgments were not influenced by exposure duration clearly did not hold for the present conditions. There was no clear evidence in our data that performance was a negatively accelerated function of exposure duration. Therefore, we tend to believe that we would not observe a leveling off by 5 seconds, and that the results of the earlier study are not compatible

with those of the present study. In retrospect, we have regretted not including a condition which used a 5-second exposure duration. Not having done so, we cannot conclude with certainty whether some disagreement between the two studies does or does not remain.

Experiment IV

It has been proposed that recency judgments are mediated primarily by a type of two-category classification learning (Underwood, 1977). The two words within a test pair are said to have been classified in a way that would produce direct temporal information. An item might be classified as falling into the first half of the list, or into the second half of the list. Or, an item might be classified as having occurred in the early part of the list, or the later part. For some pairs, the two categories may be first or second. There may be three categories, such as early, middle, late, but whatever the categories to which an item is assigned, at least one word in a test pair is learned as belonging to a category that carries direct temporal information. This theory was post hoc, and no tests of it were made. The theory simply asserts that we will understand how within-list recency judgments are made by understanding how two-category classification learning occurs. Experiments IV and V are concerned with this theory.

One of the obvious implications of the theory is that there should be an appreciable correlation between two-category classification learning and recency judgments. As discussed in detail elsewhere (Underwood, 1975), to find a strong correlation in a case like this cannot be used to support the theory, but such a correlation justifies further work involving more direct tests of the theory. If, however, the correlation is about zero,

the theory should be discarded at once. Our initial step was to determine the relationship across subjects between two-category classification learning and rate of acquisition of recency judgments. We used a two-category classification task and also one that might be described as involving three categories. For recency judgments, we used the list learned by the subjects in Experiment 1, giving five study-test trials, just as was done in Condition RR of Experiment I.

Method

Lists. Two lists (Lists 1 and 2) of 60 items each were constructed for two-category classification learning. The 120 words all had four letters, and the words were assigned randomly to one of two lists. Each list was given a single study and test trial. On the study trial, 30 of the words were underlined, 30 were not. The subjects were requested to learn which items were underlined and which were not. On the unpaced test trial, the 60 words were all shown without underlining, and the subjects were required to make a YES-NO decision for each word to indicate whether it had or had not been underlined on the study trial.

Lists 3 and 4 were also made up of four-letter words. For the study list, 40 words were presented, 20 of them underlined, 20 not. After a single study trial, the subjects were given an unpaced test consisting of 40 words from the study list plus 20 new words. The subjects were asked to indicate the appropriate classification for each item: YES, indicating the item was in the study list and was underlined; NO, that the item was in the study list but was not underlined, and NISL, that the item was not in the study list. We will speak of this as representing the classification of

words into one of three categories.

Procedure and subjects. The 30 subjects served in two sessions, these two sessions occurring on two consecutive days. On the first day, the subjects were given the four classification lists, all subjects receiving these lists in the order 1, 2, 3, and 4. On the study trial, each word was shown for 2 seconds. As already noted, there was a single study trial followed by an unpaced test trial. On the test sheet for each list, the order of the test words differed from the order used on the study trial. The recency learning was given on the second day. All procedures for recency learning were exactly the same as for Condition RR in Experiment I.

Results and Discussion

The two-category classification tests were scored to indicate the total errors. For Lists 1 and 2, the total consisted of the misses plus the false alarms on those words that were not underlined on the study trial. The means and standard deviations were 12.97 (5.44) for List 1, 9.50 (6.47) for List 2. The difference between the two means indicates learning-to-learn and was reliable (t = 3.30). The totals for List 3 and List 4 consisted of the misses on items which were underlined, false alarms on old items that were not underlined, and false alarms on new items. The means were 23.56 (7.67) and 21.63 (7.85) for Lists 3 and 4 in order. Because the total number of errors possible was equivalent for all four lists, it is evident that classifying words into one of three categories (Lists 3 and 4) was more difficult than classifying them into two categories (Lists 1 and 2).

The correlation between the scores on Lists 1 and 2 was .56; that between Lists 3 and 4 was .60. Furthermore, the two types of lists did not

correlate highly; the combined scores for Lists 1 and 2 and those for Lists 3 and 4 produced a correlation of .48. The reliability of the recency scores was determined by the correlation between trials 1, 3, and 5 combined, with that for trials 2 and 4 combined. The value was .85.

The critical correlations are those between the scores on the classification tasks and those on the recency task. The correlation between the combined scores on Lists 1 and 2 and the number of correct recency judgments was .69. The value relating Lists 3 and 4 and recency scores was .35. Thus, the data indicate a strong relationship between the two-category task (Lists 1 and 2) and recency judgments ($p < .01$). For three categories, the correlation would not be judged to be significantly different from zero, although the two correlations (.69 and .35) were not reliably different.

We will conclude that it is not unreasonable to suppose that the two-category classification task represents a paradigm that might also be involved in acquiring recency judgments. Clearly, the scores on the two-category classification tasks correlated highly with the scores on the recency-learning task. Just why the classification into three categories differed from the classification into the two categories is not known, and will not become a matter for further investigation here.

Experiment V

The major purpose of this experiment was to test the idea that recency judgments are based on a form of two-category classification learning. In this experiment, the subjects first learned to make recency judgments. Then, on a transfer test, the recency lists were changed into a two-category classification list in which the correct items on the recency test became the under-

lined items on the classification task, and the incorrect items on the recency test became nonunderlined items on the classification task. If the theory is correct, transfer between the two tasks should be heavy and positive. A control baseline was established by groups having an irrelevant recency task prior to learning the two-category classification task.

We have seen that in three experiments (I, II, III) lag length was directly related to recency learning. In an earlier study (Underwood, 1977, Experiment 11) we found no effect of lag on recency judgments. We have been unable to offer a reasonable hypothesis for the contradictions across experiments. One quite remote possibility was differences in word frequency between the earlier experiment and the present ones. Nevertheless, because we wanted to study the influence of word frequency anyhow, we have included this in the present study. There is evidence (Zimmerman, Shaughnessy, & Underwood, 1972) that associations among words in a two-category classification task does not influence the rate of acquisition. A variation in word frequency is an indirect way of manipulating the number of interitem associations among words within a list. Insofar as a word-frequency effect is dependent upon interitem associations, and insofar as recency judgments are based on two-category classification learning, we would not expect word frequency to have an influence on either the recency learning or the two-category classification task.

Method

Lists. There were four frequency levels, each represented by 60 words. In order to keep the lists as homogeneous as possible on factors other than frequency, we placed further restrictions on the words. First, all words

had four letters. Second, all words were nouns, although this did not mean that they could not function as verbs. Third, the nouns were as concrete as was possible (according to our judgment). The four frequency levels were defined by Thorndike-Lorge (1944) as follows: FL-1 was made up of words having frequencies of from 1 through 5 (G count); in FL-2 the words had frequencies of from 10 through 15; for FL-3, the frequencies ranged from 20 through 49, and for FL-4, all words were AA.

The words were placed randomly in the lists, and the pairs were chosen for the recency tests so that three lag levels were defined, with 10 pairs representing each lag. The lag levels were short (2, 3, or 4), medium (10, 11), or long (30). All 60 words in the lists were tested, i.e., there were no primacy or recency buffers.

The two-category classification lists were constructed from the same words used to form FL-1 and FL-4. Thus, the classification tasks were represented only by the two extremes of frequency. For these lists, the 30 words which were correct in the recency pairings were underlined on the study trials, and those that were incorrect were not underlined.

Procedure and subjects. There were four groups of 20 subjects each, assigned to conditions by a block-randomized schedule. The initial conditions consisted of the recency learning as a function of four frequency levels, one group being assigned to each level. There were four study-test cycles, with the words being presented at a 3-second rate on the study trials. The tests for recency discrimination were unpaced. The order of the pairs on the test sheets differed for each of the four test trials.

We will identify each group in terms of the frequency level of the

first list (recency-judgment list). Groups FL-1 and FL-2 were both given the two-category classification task using the FL-1 words used for recency judgments; Groups FL-3 and FL-4 were both given the two-category classification task using the FL-4 words. Thus, Groups FL-1 and FL-4 should show positive transfer to the classification tasks from the recency lists if recency discrimination is a case of two-category learning. Group FL-2 served as a control baseline for Group FL-1 on the two-category task, and Group FL-3 served as the control for Group FL-4. Finally, a direct comparison between Groups FL-3 and FL-2 on the two-category classification task would tell us if extreme differences in word frequency influenced learning.

All subjects were given two study-test cycles on the two-category classification lists. The study trials were conducted at a 3-second rate, and the test trials were paced, using the same rate. The subjects responded YES or NO to each item to indicate whether it had or had not been underlined on the test trials. It seemed necessary to use the paced test to prevent the subjects from reconstructing the recency list and then deriving classification information from the reconstructed list. The subjects were fully informed concerning the nature of the two-category classification test prior to the first study trial and again before the first test trial. The subjects in Groups FL-1 and FL-4 were not told that the underlined words in the classification task were the correct words for the recency pairs. A different order of the words was used on each of the four trials (two study and two test) for the two-category classification task.

Results

Recency judgments. Word frequency did not have a reliable influence

on the recency judgments. The mean correct recency judgments per trial (with standard deviations) for Groups FL-1 through FL-4 were 6.98 (.75), 6.89 (.81), 6.90 (.99), and 6.27 (1.06). Because the means are based on 10 recency tests, the values can be changed to percentages by moving the decimal point one place to the right. Both lag (F = 75.82) and trials (F = 35.51) produced reliable effects, and the interaction between the two can be considered reliable, F (6,456) = 2.50, $p < .05$. The results for these two variables are plotted in Figure 8. As may be seen, the statistical interaction between the two variables does not appear to be one with systematic importance.

Two-category classification. It will be remembered that on the two-category classification task, Group FL-2 served as the control for Group FL-1, and Group FL-3 served as the control for Group FL-4. The mean performance measures are shown in Figure 9, with the scores for the two groups tested on low-frequency words being in the left panel, the scores for the two groups tested on high-frequency words being in the right panel. Since there were 60 items on the classification task, the scores range roughly from 67% correct to 88% correct. A correct response was identified as responding YES when a word had been underlined and NO when it had not been underlined on the study trial. The surprising fact shown by these data is that the transfer was negative; performance on the classification task was retarded by having it made up of words which had previously appeared in the recency-discrimination task. This was true in spite of the fact that all

underlined words on the classification task had been correct words for the recency discriminations. The negative transfer effect was reliable, F (1,76) = 5.72, $p < .02$, and although the amount of negative transfer appears less for the low-frequency words (left panel) than for the high-frequency words (right panel), the interaction was not reliable, F (1,76) = 2.08, $p > .05$. In any absolute sense, the negative transfer was not great, but that no positive transfer was observed is obviously very troublesome for the theory which presumes that recency discriminations represent a form of two-category classification learning.

One other fact given in Figure 9 must be mentioned. The difference between Groups FL-2 and FL-3 is small and nonsignificant statistically. The two-category classification task given Group FL-2 was made up of low-frequency words, the task given Group FL-3 was made up of high-frequency words. The implication is that word frequency does not influence two-category classification learning.

Correlations. The theory that recency discrimination fits the paradigm of the two-category classification task leads to the expectation that the correlation between recency scores and classification scores should be very high. The correlation for Group FL-1 was .29, that for Group FL-4, .55. The theory could probably tolerate the second correlation, but certainly not the first. Experiment IV showed that the skills involved in learning recency judgments were rather highly correlated with the skills involved in classification learning. Groups FL-2 and FL-3 provide further

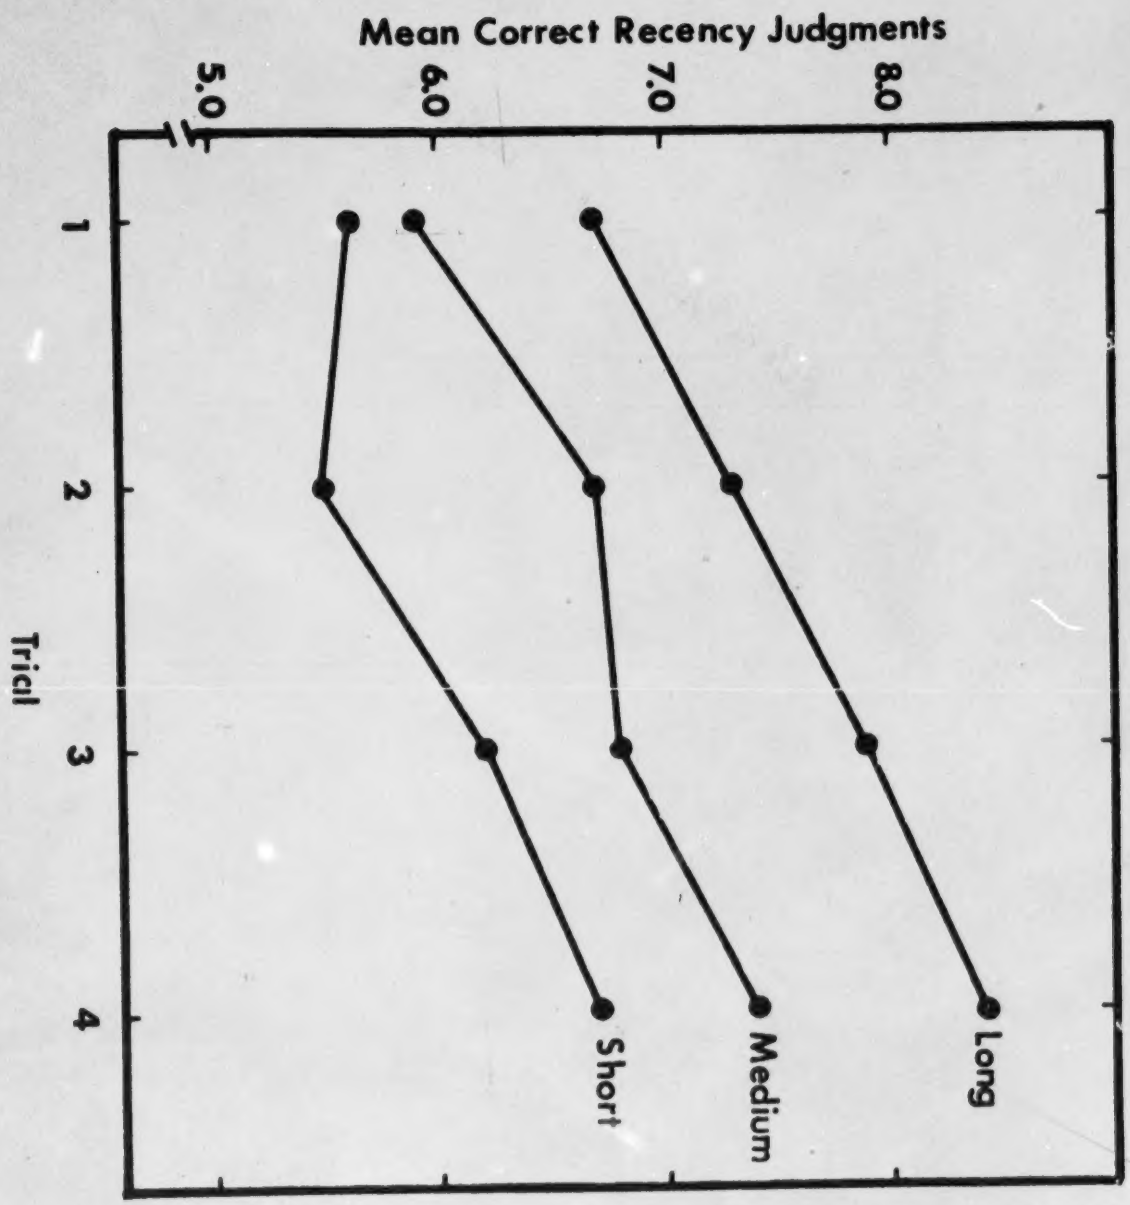

Figure 8. Recency judgments as a function of lag and trial. Experiment V.

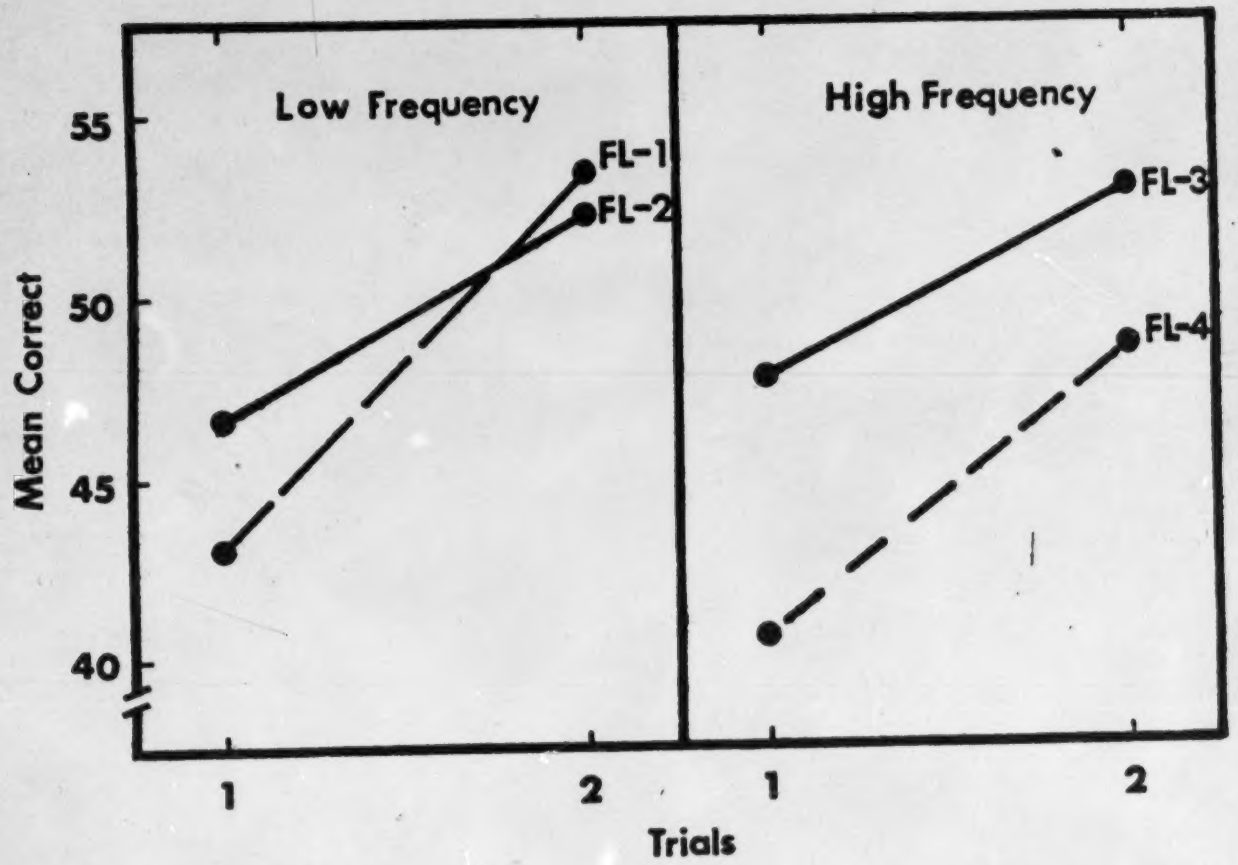

Figure 9. Transfer from recency judgments to two-category classification learning of the same words as related to frequency level and trial on the classification learning. Experiment V.

tests of this relationship. For FL-2, the correlation was .63, for FL-3, .51.

Discussion

Two findings seriously question the idea that the learning of recency discriminations is to be understood by studying the learning underlying the two-category classification task. The first finding is the failure to find positive transfer from recency learning to two-category learning. Even if several different categories were used to classify items in learning recency discriminations, it seems to us that the category label representing the most recent item in a pair could be translated readily into the appropriate category (underlining) on the two-category classification task. Strangely enough, the fact that there was negative transfer might be viewed as being more favorable toward the theory than would be true had there been no transfer of any kind. The negative transfer indicates that the two tasks made contact in some way; the associations involved in both tasks had some overlap. Nevertheless, the nature of the overlap does not follow that expected by our thinking.

The second finding which is difficult for our theory is the relatively low correlation between the scores for the two types of learning for Group FL-1. It must necessarily follow from the theory that a subject who learns recency discriminations rapidly must also learn the two-category classification task based on the same items rapidly, and this was not found.

Had the basic evidence been in support of our theory, we could point also to the fact that word frequency had no influence on either recency discriminations or on two-category classification learning as supporting the idea that the two tasks have high commonality. But, without direct support

for the theory, the lack of an effect of word frequency on the two tasks carries little theoretical weight.

Experiment VI

We have seen that the acquisition of recency discriminations between two words within a relatively long list is probably not a special case of learning a two-category classification. In the present experiment, we will make a test of a second theoretical notion proposed in an earlier report (Underwood, 1977). It was assumed that temporal discriminations between and within lists are primarily derived from associative learning. The idea that within-list recency discriminations result primarily from two-category classification learning was an attempt to be more specific about the nature of the associative learning involved. To say that temporal codes are primarily mediated by associative learning indicates that a second factor is involved. We have called this second factor the recency principle. This principle simply states that immediately after the presentation of an item, the temporal discrimination between it and all items preceding it is perfect, but that as time passes, the temporal discrimination becomes less and less reliable.

The recency principle presumes that information for temporal discrimination is given directly in the memory and is not based on associative learning. In this sense it is a primitive mechanism which establishes some degree of orderliness in memory functioning. For the principle to be maximally useful, the rate at which frequency information is lost should be specified. We believe that very short intervals are involved in that the event which is most recent (in fact, and in memory) will normally be replaced by another event within a few seconds. However, this loss is probably

influenced by the nature of the activity which occurs after the appearance of the critical event. Thus, the loss of recency may occur within a few seconds, and may be due to a fading with time, to a fading produced by other events, or to both.

We do not think there is a perfectly appropriate way to produce empirical evidence for the gradient describing the loss in recency over time. The basic idea would be to try to show that if the subject is shown a series of items ending with X, the recency discrimination between X and the previous items would be perfect at a zero-retention interval, but would fall as the interval increases. The problem arises in deciding what the subject should be required to do during the retention interval, and we could find no solution which appeared entirely satisfactory to us. In the end, we chose not to fill the interval with extraneous activity, but at the same time we arranged the tests so that the subject could not anticipate tests on particular items.

Method

Basic design. The lists all consisted of 10 words, and the central interest was directed at the recency judgments in which the 10th word was involved. The memory drum was set for a 3-second rate. Following the appearance of the 10th word, the drum continued to turn and the events were as follows for the successive positions: 11th, recency test; 12th, blank; 13, recency test; 14th and 15th, blank; 16th, recency test; 17th, 18th, 19th, blank; 20th, recency test; 21st, blank; 22nd, the word STUDY appeared, and on the next exposure the first word from another list of 10 words was shown. As can be seen, four recency tests were given for each list, hence, 8 of the 10 words in each list were tested. The four recency tests occurred 0, 6, 15, and 27 seconds after the appearance of the 10th word. One of these

tests always involved the 10th word, and across lists the 10th word was tested equally often after each of the four retention intervals. Our major interest, of course, was in the change in the number of correct recency judgments as the retention interval became longer and longer. The recency principle assumes that with the zero interval performance will be essentially perfect. The critical question concerns the rate of fall in the number of correct recency judgments over time.

Lists. Each subject was given 64 lists, and since the 10th item was always tested for each list, there were 16 tests at each of the four retention intervals. On the recency tests the 10th item was paired with the 9th, 8th, 6th, and 3rd items, each for an equal number of times. Thus, the lag and the retnetion interval were orthogonal, and produced 16 different conditions. However, there were only four different items in each of the 16 conditions and our intent was not to be concerned with the influence of lag variable.

All of the words contained four letters and were assigned to lists and to positions within the lists on a random basis. After the two items for the critical test had been selected for each list, three additional recency tests were devised by drawing three additional pairs randomly from among the eight remaining. These pairs were in turn assigned randomly to the three remaining retention intervals.

Procedure and subjects. The experiment involved a single group of 20 subjects, and all were tested by a single experimenter. We had anticipated that subjects might have difficulty responding on the recency tests within the 3-second rate period. Therefore, subjects were given practice on three

lists. We were prepared to give additional practice lists if needed but this was found to be unnecessary. The subjects were given 32 experimental lists on the first day and then returned the following day for completion of the remaining lists. A short rest period was given after each block of four lists. The order in which the lists was given was the same for all 20 subjects.

Results

The recency judgments involving the 10th item at each retention interval are shown in Figure 10, expressed as percent correct. Nine of the 20 subjects failed to perform perfectly on all 16 tests with the zero-retention interval, but performance was near perfect for the group (96%). After a 6-second retention interval, there was a drop of approximately 30%, but no further drop of consequence occurred as the interval lengthened beyond 6 seconds. Overall, the differences among the four intervals was reliable, $F\ (3,57) = 36.23$, $p < .01$. We interpret these data to support the basic idea of the recency principle. Furthermore, the data suggest that when responding based on the recency principle is lost, it is lost completely; performance does not change over the longer retention intervals. We do not know, of course, whether the loss was entirely time dependent or whether it was produced by the recency test given at the zero-retention interval, or by some combination of the two.

It was suggested elsewhere (Underwood, 1977) that the recency principle was probably not all or none, meaning that an item that was next to the last

item in the list might show a recency effect. Thus, in the present experiment, this would mean that recency tests at the zero-retention interval would be more frequently correct when the 9th item and an earlier one constituted the test pair than when the 8th item plus an earlier one constituted the test pair. We have made a test of this. We examined all recency tests in which the 9th item was paired with one from positions 1 through 8 (there were 11 such cases), those in which the 8th item was paired with one from positions 1 through 7 (9 cases), and those in which the 7th item was paired with one from positions 1 through 6 (12 cases), and for which, in all cases, the zero-retention interval was involved, i.e., the recency test was given immediately after the presentation of the 10th word. The percent correct recency judgments for these three cases, plus those for the words in position 10, are shown in Figure 11. The baseline shows the study position of the correct word. There is certainly no evidence that a recency-like effect held for items that occurred prior to the 10th item in the study lists. Apparently, the occurrence of the 10th item in the study list destroys recency information for the 9th item.

It should be mentioned that the above represents a strong test of the role of recency for other than the last item. The mean lags between the two items in the test pairs necessarily decreased a little as the position of the correct item moved further and further back in the list (from 9, to 8, to 7). This could have decreased performance correspondingly, resulting in a gradient extending across all four points in Figure 11. This obviously

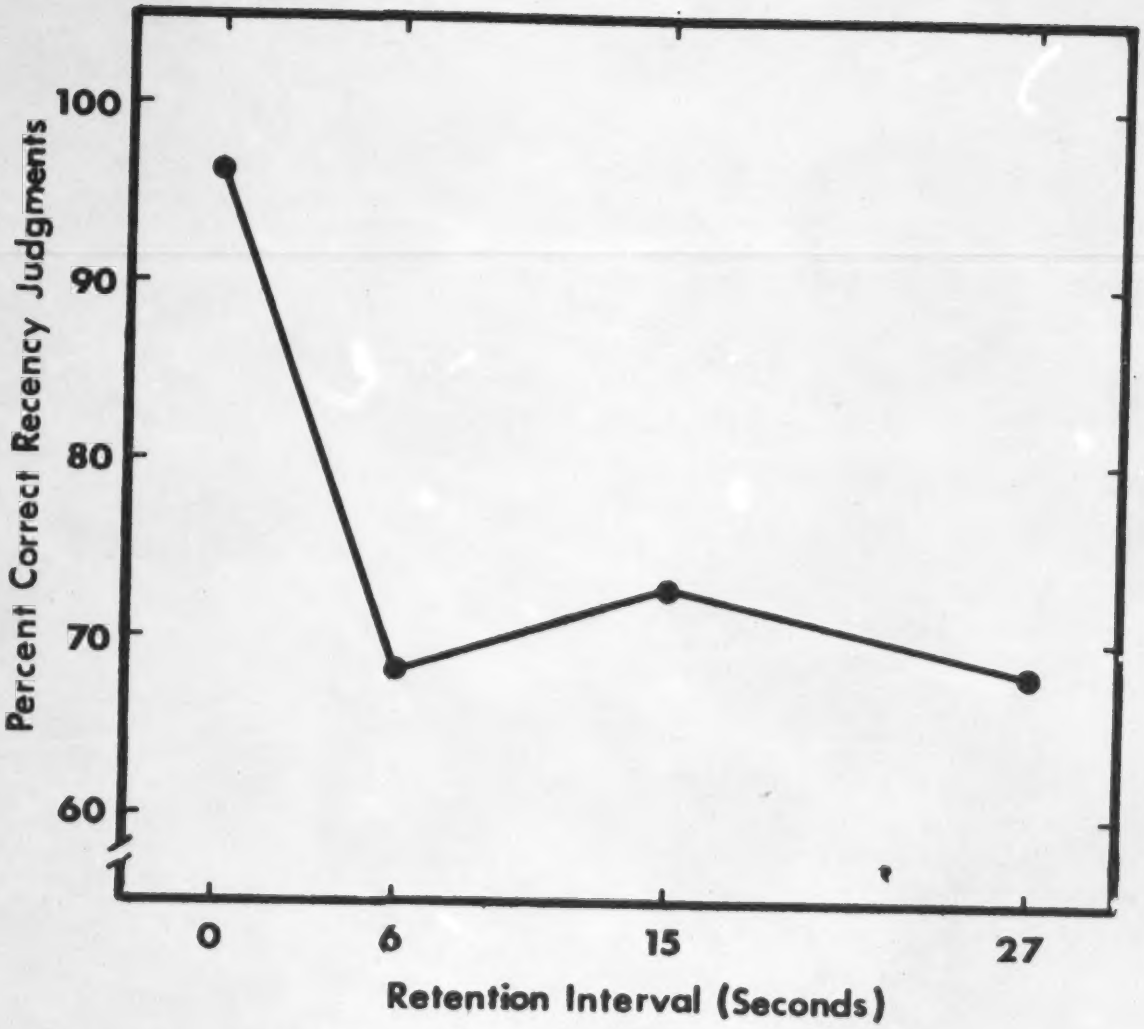

Figure 10. Recency judgments as a function of the retention interval when the correct item was the last word in a 1-word list. Experiment VI.

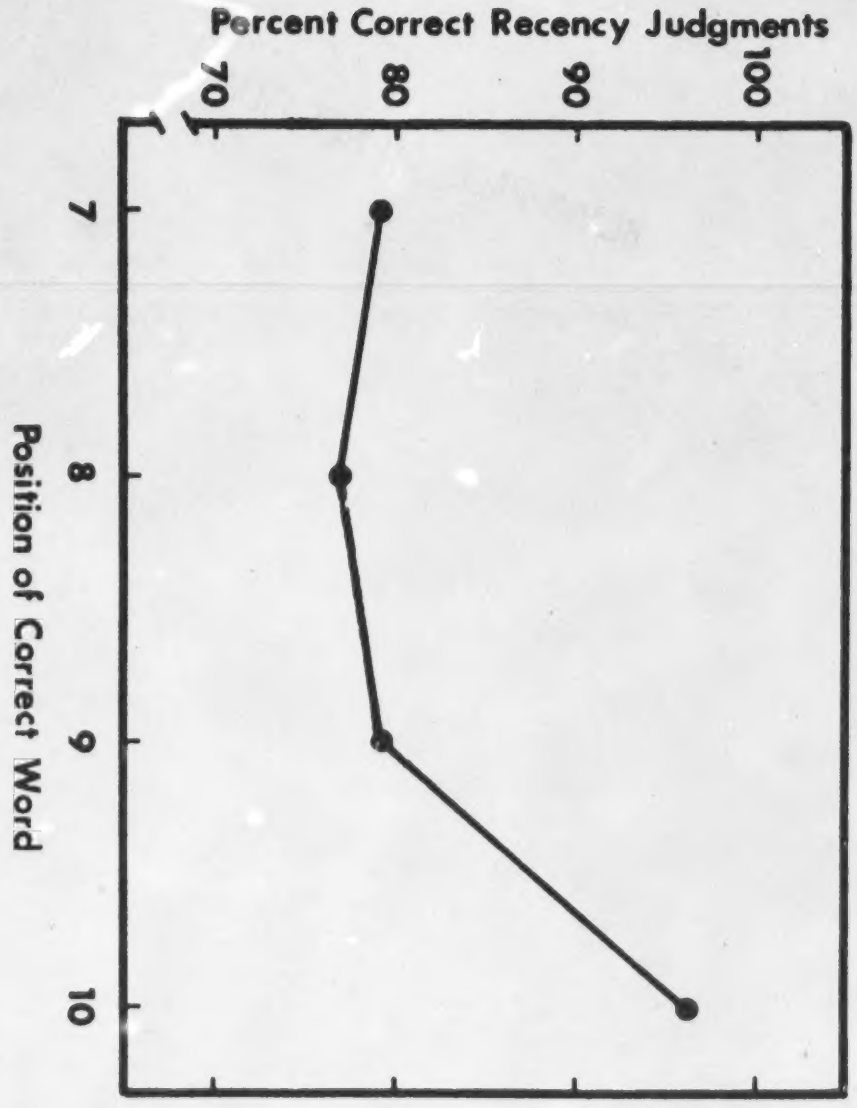

Figure 11. Recency judgments as related to the study position of the correct word. Experiment VI.

did not happen.

We have examined the recency judgments on all of the noncritical items, but the only fact that might seem somewhat surprising was the complete lack of any effect of the length of the retention interval. There was simply no relationship. For example, we may give the values when the item in position 7 of the study list was the correct (most recent) item. The correct recency judgments were 80, 73, 73, and 78% for the four retention intervals in order.

Discussion

We conclude that this experiment shows the presence of a short-term component involved in recency judgments, a component which is consistent with the basic idea of the recency principle. We emphasize that our results do not determine whether the loss of recency of the short term kind is due to being tested, or whether it is entirely time dependent. If it is entirely time dependent, the time span over which it operates is shorter than we had originally contenplated. There was simply no influence produced by the retention interval beyond 6 seconds regardless of the position in the study list held by the correct item. It may be noted that Brown (1973), using pictures as the stimuli, found performance on recency judgments to decrease as the retention interval increased, but this was with a procedure in which the intervals were filled with new items which the subjects were to remember.

The evidence has suggested that the recency principle as it was manifested in our data was something of an all-or-none affair. All loss of recency occurred within 6 seconds, and there was no further loss beyond that point. Unless the item was last in the study list, no loss could be attributed to a recency effect. Thus, an item in position nine in the study list

showed no evidence of a loss over time. We rather suspect that the all-or-none like characteristics will not always be found.

General Discussion

The data have shown that lag judgments are very poor response measures for indexing temporal coding. Although our subjects did show a small amount of learning reflecting differences in lag lengths, only the magnitude of the judgments for long lags increased appropriately; the magnitude of the judgments for the short lags did not decrease as they should have if the subjects recognized that their earlier judgments had been too long. In spite of the poor performance shown by the subjects when making lag judgments, they acquired a great deal of knowledge about the temporal relationship of the items in the list. This was shown by the heavy positive transfer when subjects were switched from lag judgments to either recency or position judgments. In some cases the transfer was essentially complete, i.e., while the subjects were struggling with lag judgments they learned as much about the order of the items in the lists (as measured by recency judgments) as did the subjects who made recency judgments at all times. As a performance measure, the lag judgment simply does not reflect the information which the subjects have acquired about the temporal relationships of the items within the list.

We had assumed that recency judgments are mediated by two factors. One of these factors consists of associative learning. In some cases this is serial learning, although we did not study this matter directly in the present experiments. In other cases, which are probably more typical, we assumed that recency judgments result from a learning to categorize each item into one of a limited number of categories specifying a temporal dimen-

sion directly or indirectly. Usually the dimension is represented by only two categories, such as early and late. We had been led to this idea in part because of an earlier study (Underwood, 1977) in which it was found that lag length did not relate to the number of correct recency judgments. Four of the present experiments showed that lag length was related to the number of correct recency judgments, and we have not discovered why the earlier experiment should have differed on this matter. In any event, our notion was that the associative learning involved in acquiring temporal codes for items within a list was a form of two-category classification learning. This possibility was heightened by the fact that performance in learning recency judgments correlated with performance in learning a two-category classification task. However, when (Experiment V) we made a direct test of the theory, it was found to be wanting. We must conclude that the evidence does not allow us to maintain the idea that the associative learning in within-list temporal coding can be described in the same way we would describe the learning which occurs when the subjects learn to classify events into one of two categories.

The second factor which we believe must be considered in trying to understand temporal coding is a very short term factor which we have said represents a recency principle. The evidence from Experiment VI makes it reasonable to conclude this. Nevertheless, it will be necessary in the long run to assess the degree to which the gradient resulting from the recency principle is purely time dependent and the degree to which it can be changed by other activites.

As explained elsewhere (Underwood, 1977), our original interest in

temporal coding arose because of the role we found it playing in proactive inhibition between lists. As frequently happens, the interest in the issues involved in temporal coding per se has come to be self-sustaining without referring them to proactive inhibition.

References

Battig, W. F., & Montague, W. E. Category norms for verbal items in 56 categories: A replication and extension of the Connecticut norms. Journal of Experimental Psychology Monograph, 1969, 80 (3, Pt 2).

Brown, A. L. Judgments of recency for long sequences of pictures: The absence of a developmental trend. Journal of Experimental Child Psychology, 1973, 15, 473-480.

Hintzman, D. L., & Block, R. A. Memory for the spacing of repetitions. Journal of Experimental Psychology, 1973, 99, 70-[illegible].

Hintzman, D. L., Summers, J. J., & Block, R. A. Spacing judgments as an index of study-phase retrieval. Journal of Experimental Psychology: Human Learning and Memory, 1975, 1, 31-40.

Thorndike, E. L, & Lorge, I. The teacher's wordbook of 30,000 words. New York: Bureau of Publications, Teachers College, 1944.

Toglia, M. P., & Kimble, G. A. Recall and use of serial position information. Journal of Experimental Psychology: Human Learning and Memory, 1976, 2, 431-445.

Underwood, B. J. Individual differences as a crucible in theory construction. American Psychologist, 1975, 30, 128-134.

Underwood, B. J. Temporal codes for memories: Issues and problems. Hillsdale, New Jersey: L. Erlbaum Associates, 1977.

Zimmerman, J., Shaughnessy, J. J., & Underwood, B. J. The role of associations in verbal-discrimination learning. American Journal of Psychology, 1972, 85, 499-518.